










Engineering Innovations

Propulsion	
Thermal Protection Systems	
Materials and Manufacturing	
Aerodynamics and Flight Dynamics	
Avionics, Navigation, and Instrumentation	
Software	
Structural Design	
Robotics and Automation	
Systems Engineering for Life Cycle of Complex Systems	



Propulsion

Introduction

Yolanda Harris

Space Shuttle Main Engine

Fred Jue

Chemochromic Hydrogen

Leak Detectors

Luke Roberson

Janine Captain

Martha Williams

Mary Whitten

The First Human-Rated

Reusable Solid Rocket Motor

Fred Perkins

Holly Lamb

Orbital Propulsion Systems

Cecil Gibson

Willard Castner

Robert Cort

Samuel Jones

Pioneering Inspection Tool

Mike Lingbloom

Propulsion Systems and

Hazardous Gas Detection

Bill Helms

David Collins

Ozzie Fish

Richard Mizell

The launch of the Space Shuttle was probably the most visible event of the entire mission cycle. The image of the Main Propulsion System—the Space Shuttle Main Engine and the Solid Rocket Boosters (SRBs)—powering the Orbiter into space captured the attention and the imagination of people around the globe. Even by 2010 standards, these main engines' performance was unsurpassed compared to any other engines. They were a quantum leap from previous rocket engines. The main engines were the most reliable and extensively tested rocket engine before and during the shuttle era.

The shuttle's SRBs were the largest ever used, the first reusable rocket, and the only solid fuel certified for human spaceflight. This technology, engineering, and manufacturing may remain unsurpassed for decades to come.


But the shuttle's propulsion capabilities also encompassed the Orbiter's equally important array of rockets—the Orbital Maneuvering System and the Reaction Control System—which were used to fine-tune orbits and perform the delicate adjustments needed to dock the Orbiter with the International Space Station. The design and maintenance of the first reusable space vehicle—the Orbiter—presented a unique set of challenges. In fact, the Space Shuttle Program developed the world's most extensive materials database for propulsion. In all, the shuttle's propulsion systems achieved unprecedented engineering milestones and launched a 30-year era of American space exploration.

Space Shuttle Main Engine

NASA faced a unique challenge at the beginning of the Space Shuttle Program: to design and fly a human-rated reusable liquid propulsion rocket engine to launch the shuttle. It was the first and only liquid-fueled rocket engine to be reused from one mission to the next during the shuttle era. The improvement of the Space Shuttle Main Engine (SSME) was a continuous undertaking, with the objectives being to increase safety, reliability, and operational margins; reduce maintenance; and improve the life of the engine's high-pressure turbopumps.

The reusable SSME was a staged combustion cycle engine. Using a mixture of liquid oxygen and liquid hydrogen, the main engine could attain a maximum thrust level (in vacuum) of 232,375 kg (512,300 pounds), which is equivalent to greater than 12,000,000 horsepower (hp). The engine also featured high-performance fuel and oxidizer turbopumps that developed 69,000 hp and 25,000 hp, respectively. Ultra-high-pressure operation of the pumps and combustion chamber allowed expansion of hot gases through the exhaust nozzle to achieve efficiencies never previously attained in a rocket engine.

Requirements established for Space Shuttle design and development began in the mid 1960s. These requirements called for a two-stage-to-orbit vehicle configuration with liquid oxygen (oxidizer) and liquid hydrogen (fuel) for the Orbiter's main engines. By 1969, NASA awarded advanced engine studies to three contractor firms to further define designs necessary to meet the leap in performance demanded



by the new Space Transportation System (STS).

In 1971, the Rocketdyne division of Rockwell International was awarded a contract to design, develop, and produce the main engine.

The main engine would be the first production-staged combustion cycle engine for the United States. (The Soviet Union had previously demonstrated the viability of staged combustion cycle in the Proton vehicle in 1965.) The staged combustion cycle yielded high efficiency in a technologically advanced and complex engine that operated at pressures beyond known experience.

The design team chose a dual-preburner powerhead configuration to provide precise mixture ratio and throttling control. A low- and high-pressure turbopump, placed in series for each of the liquid hydrogen and liquid oxygen loops, generated high pressures across a wide range of power levels.

A weight target of 2,857 kg (6,300 pounds) and tight Orbiter ascent envelope requirements yielded a compact design capable of generating a nominal chamber pressure of 211 kg/cm² (3,000 pounds/in²)—about four times that of the Apollo/Saturn J-2 engine.



Michael Coats

Pilot on STS-41D (1984).
Commander on STS-29 (1989)
and STS-39 (1991).



A Balky Hydrogen Valve Halts Discovery Liftoff

"I had the privilege of being the pilot on the maiden flight of the Orbiter Discovery, a hugely successful mission. We deployed three large communications satellites and tested the dynamic response characteristics of an extendable solar array wing, which was a precursor to the much-larger solar array wings on the International Space Station.

"But the first launch attempt did not go quite as we expected. Our pulses were racing as the three main engines sequentially began to roar to life, but as we rocked forward on the launch pad it suddenly got deathly quiet and all motion stopped abruptly. With the seagulls screaming in protest outside our windows, it dawned on us we weren't going into space that day. The first comment came from Mission Specialist Steve Hawley, who broke the stunned silence by calmly saying 'I thought we'd be a lot higher at MECO (main engine cutoff).' So we soon started cracking lousy jokes while waiting for the ground crew to return to the pad and open the hatch. The joking was short-lived when we realized there was a residual fire coming up the left side of the Orbiter, fed from the same balky hydrogen valve that had caused the abort. The Launch Control Center team was quick to identify the problem and initiated the water deluge system designed for just such a contingency. We had to exit the pad elevator through a virtual wall of water. We wore thin, blue cotton flight suits back then and were soaked to the bone as we entered the air-conditioned astronaut van for the ride back to crew quarters. Our drenched crew shivered and huddled together as we watched the Discovery recede through the rear window of the van, and as Mike Mullane wryly observed, 'This isn't exactly what I expected spaceflight to be like.' The entire crew, including Commander Henry Hartsfield, the other Mission Specialists Mike Mullane and Judy Resnik, and Payload Specialist Charlie Walker, contributed to an easy camaraderie that made the long hours of training for the mission truly enjoyable."

For the first time in a boost-to-orbit rocket engine application, an on-board digital main engine controller continuously monitored and controlled all engine functions. The controller initiated and monitored engine parameters and adjusted control valves to maintain the performance parameters required by the mission. When detecting a malfunction, it also commanded the engine into a safe lockup mode or engine shutdown.

Design Challenges

Emphasis on fatigue capability, strength, ease of assembly and disassembly, maintainability, and materials compatibility were all major considerations in achieving a fully reusable design.

Specialized materials needed to be incorporated into the design to meet the severe operating environments. NASA successfully adapted advanced alloys, including cast titanium, Inconel® 718 (a high-strength, nickel-based superalloy used in the main combustion chamber support jacket and powerhead), and NARloy-Z (a high-conductivity, copper-based alloy used as the liner in the main combustion chamber). NASA also oversaw the development of single-crystal turbine blades for the high-pressure turbopumps. This innovation essentially eliminated the grain boundary separation failure mechanism (blade cracking) that had limited the service life of the pumps. Nonmetallic materials such as Kel-F® (a plastic used in turbopump seals), Armalon® fabric (turbopump bearing cage material), and P5N carbon-graphite seal material were also incorporated into the design.

Material sensitivity to oxygen environment was a major concern for compatibility due to reaction and



ignition under the high pressures. Mechanical impact testing had vastly expanded in the 1970s to accommodate the shuttle engine's varied operating conditions. This led to a new class of liquid oxygen reaction testing up to 703 kg/cm² (10,000 pounds/in²).


Engineers also needed to understand long-term reaction to hydrogen effects to achieve full reusability. Thus, a whole field of materials testing evolved to evaluate the behavior of hydrogen charging on all affected materials.

NASA developed new tools to accomplish design advancements. Engineering design tools advanced along with the digital age as analysis migrated from the mainframe platform to workstations and desktop personal computers. Fracture mechanics and fracture control became critical tools for understanding the characteristics of crack propagation to ensure design reusability. As the analytical tools and processor power improved over the decades, cycle time for engineering analysis such as finite element models, computer-aided design and manufacturing, and computational fluid dynamics dropped from days to minutes. Real-time engine performance analyses were conducted during ground tests and flights at the end of the shuttle era.

Development and Certification

The shuttle propulsion system was the most critical system during ascent; therefore, a high level of testing was needed prior to first flight to demonstrate engine maturity. Component-level testing of the preburners and thrust chamber began in 1974 at Rocketdyne's Santa Susana Field Laboratory in Southern California.

The first engine-level test of the main engine—the Integrated Subsystem



A 1970s-era Space Shuttle Main Engine undergoes testing at Rocketdyne's Santa Susana Field Laboratory near Los Angeles, California.

Test Bed—occurred in 1975 at the NASA National Space Technology Laboratory (now Stennis Space Center) in Mississippi and relied on facility controls, as the main engine controller was not yet available.

NASA and Rocketdyne pursued an aggressive test schedule at their respective facilities. Stennis Space Center with three test stands and Rocketdyne with one test stand completed 152 engine tests in 1980 alone—a record that has not been exceeded since. This ramp-up to 100,000 seconds represented a team effort of personnel and facilities to overachieve a stated development goal of 65,000 seconds set by then-Administrator John Yardley as the maturity level deemed flightworthy. NASA verified operation at altitude conditions and also demonstrated the rigors of sea-level performance and engine gimbaling for thrust vector control. The Rocketdyne laboratory supplemented sea-level testing as well as deep throttling by using a low 33:1 expansion ratio nozzle. This testing was crucial in identifying shortcomings

related to the initial design of the high-pressure turbopumps, powerhead, valves, and nozzles.

Extensive margin testing beyond the normal flight envelope—including high-power, extended-duration tests and near-depleted inlet propellant conditions to simulate the effects of microgravity—provided further confidence in the design. Engineers subjected key components to a full series of design verification tests, some with intentional hardware defects, to validate safety margins should the components develop undetected flaws during operation.

NASA and Rocketdyne also performed system testing to replicate the three engine cluster interactions with the Orbiter. The Main Propulsion Test Article consisted of an Orbiter aft fuselage, complete with full thrust structure, main propulsion electrical and system plumbing, External Tank, and three main engines. To validate that the Main Propulsion System was ready for launch, engineers completed 18 tests at the National Space Technology Laboratory by 1981.



The completion of the main engine preliminary flight certification in March 1981 marked a major milestone in clearing the initial flights at 100% rated power level.

Design Evolutions

A major requirement in engine design was the ability to operate at various power levels. The original engine life requirement was 100 nominal missions and 27,000 seconds (7.5 hours) of engine life. Nominal thrust, designated as rated power level, was 213,189 kg (470,000 pounds) in vacuum. The life requirement included six exposures at the emergency power level of 232,375 kg (512,300 pounds), which was designated 109% of rated power level. To maximize the number of missions possible at emergency power level, an assessment of the engine capability resulted in reducing the number of nominal missions per engine to 55 missions at 109%. Emergency power level was subsequently renamed full power level.

Ongoing ascent trajectory analysis determined 65% of rated power level to be sufficient to power the vehicle through its period of maximum aerodynamic pressure during ascent. Minimum power level was later refined upward to 67%.

On April 12, 1981, Space Shuttle Columbia lifted off Launch Pad 39A from Kennedy Space Center in Florida on its maiden voyage. The first flight configuration engines were aptly named the First Manned Orbital Flight SSMEs. These engines were flown during the initial five shuttle development missions at 100% rated power level thrust. Work done to prepare for the next flight validated the ability to perform

routine engine maintenance without removing them from the Orbiter.

The successful flight of STS-1 initiated the development of a full-power (109% rated power level) engine. The higher thrust capability was needed to support an envisioned multitude of NASA, commercial, and Department of Defense payloads, especially if the shuttle was launched from the West Coast. By 1983, however, test failures demonstrated the basic engine lacked margin to continuously operate at 109% thrust, and full-power-level development was halted. Other engine improvements were implemented into what was called the Phase II engine. During this period, the engine program was restructured into two programs—flight and development.

Post-Challenger Return to Flight

The 1986 Challenger accident provoked fundamental changes to the shuttle, including an improved main engine called Phase II. This included changes to the high-pressure turbopumps and main combustion chamber, avionics, valves, and high-pressure fuel duct insulation. An additional 90,241 seconds of engine testing accrued, including recertification to 104% rated power level.

The new Phase II engine continued to be the workhorse configuration for shuttle launches up to the late 1990s while additional improvements envisioned during the 1980s were undergoing development and flight certification for later incorporation. NASA targeted five major components for advanced development to further enhance safety and reliability, lower recurring costs, and increase performance capability. These components included the powerhead, heat exchanger, main combustion

chamber, and high-pressure oxidizer and fuel turbopumps.

These major changes would later be divided into two “Block” configuration upgrades, with Rocketdyne tasked to improve the powerhead, heat exchanger, and main combustion chamber while Pratt & Whitney was selected to design, develop, and produce the improved high-pressure turbopumps.

Pratt & Whitney Company of United Technologies began the effort in 1986 to provide alternate high-pressure turbopumps as direct line replaceable units for the main engines. Pratt & Whitney used staged combustion experience from its development of the XLR-129 engine for the US Air Force and cryogenic hydrogen experience from the RL-10 (an upper-stage engine used by NASA, the military, and commercial enterprises) along with SSME lessons learned to design the new pumps. The redesign of the components eliminated critical failure modes and increased safety margins.

Next Generation

The Block I configuration became the successor to the Phase II engine. A new Pratt & Whitney high-pressure oxygen turbopump, an improved two-duct engine powerhead, and a single-tube heat exchanger were introduced that collectively used new design and production processes to eliminate failure causes. Also it increased the inherent reliability and operating margin and reduced production cycle time and costs. This Block I engine first flew on STS-70 (1995).

The powerhead redesign was less risky and was chosen to proceed ahead of the main combustion chamber.



© Pratt & Whitney Rocketdyne. All rights reserved.

The Technology Test Bed Space Shuttle Main Engine test program was conducted at Marshall Space Flight Center, Alabama, between September 1988 and May 1996. The program demonstrated the ability of the main engine to accommodate a wide variation in safe operating ranges.

The two-duct powerhead eliminated 74 welds and had 52 fewer parts. This improved design led to production simplification and a 40% cost reduction compared to the previous three-duct configuration. The two-duct configuration provided an improvement to the hot gas flow field distribution and reductions in dynamic pressures. The improved heat exchanger eliminated all inter-propellant welds, and its wall thickness was increased by 25% for added margin against penetration by unexpected foreign debris impact.

The new high-pressure oxygen turbopump eliminated 293 welds, added improved suction performance, and introduced a stiff single-piece disk/shaft configuration and thin-cast turbine blades. The oxygen turbopump incorporated silicon nitride (ceramic) ball bearings in a rocket engine application and could be serviced without removal from the engine. Initial

component-level testing occurred at the Pratt & Whitney West Palm Beach, Florida, testing facilities. Testing then graduated to the engine level at Stennis Space Center as well as at Marshall Space Flight Center's (MSFC's) Technology Test Bed test configuration.

The large-throat main combustion chamber began prototype testing at Rocketdyne in 1988. But it was not until 1992, after a series of combustion stability tests at the MSFC Technology Test Bed facility, that concerns regarding combustion stability were put to rest. The next improved engine—Block II—incorporated the new high-pressure fuel turbopump, modified low-pressure turbopumps, software operability enhancements, and previous Block I upgrades. These upgrades were needed to support International Space Station (ISS) launches with their heavy payloads beginning in 1998.


As Block II development testing progressed, the engineering accomplishments on the large-throat main combustion chamber matured more rapidly than the high-pressure fuel turbopump.

By February 1997, NASA had decided to go forward with an interim configuration called the Block IIA. Using the existing Phase II high-pressure fuel pump, this configuration would allow early implementation of the large-throat main combustion chamber to support ISS launches. The large-throat main combustion chamber was simpler and producible. The new chamber lowered the engine's operating pressures and temperatures while increasing the engine's operational safety margin. Changes to the low-pressure turbopumps to operate in this derated environment, along with further avionics improvements, were flown in 1998 on STS-89.

The large-throat main combustion chamber became one of the most significant safety improvements for the main engine by effectively reducing operating pressures and temperatures up to 10% for all subsystems. This design also incorporated improved cooling capability for longer life and used high-strength castings, thus eliminating 50 welds.

By the time the first Block IIA flew on STS-89 in January 1998, the large-throat main combustion chamber design had accumulated in excess of 100,000 seconds of testing time. By late 1999, the Block II high-pressure fuel turbopump had progressed into certification testing. The design philosophy mirrored those proven successful in the high-pressure oxidizer turbopump and included the elimination of 387 welds

The Improved Space Shuttle Main Engine Powerhead Component Arrangement for Block II Engines



The Block II engine combined a new high-pressure fuel turbopump with the previously flown redesigned high-pressure oxygen turbopump. Risk analysis showed that the Block II engine was twice as safe as the 1990s-era engine. Beginning with STS-110 in April 2002, all shuttle flights were powered by the improved Space Shuttle Main Engine.

and incorporation of a stiff single-piece disk/shaft, thin-cast turbine blades, and a cast pump inlet that improved the suction performance and robustness against pressure surges. As with the high-pressure oxidizer turbopump, the high-pressure fuel turbopump turbine inlet did not require off-engine inspections, which contributed significantly to improving engine turnaround time. The high-pressure fuel turbopump also demonstrated that a turbine blade failure would result in a contained, safe engine shutdown. By introducing the added operational margin of the large-throat main combustion chamber with the new turbopumps, quantitative risk analysis

projected that the Block II engine was twice as safe as the Phase II engine.

The first two single-engine flights of Block II occurred on STS-104 and STS-108 in July 2001 and December 2001, respectively, followed by the first three-engine cluster flight on STS-110 in April 2002. The high-pressure fuel turbopump had accumulated 150,843 seconds of engine test maturity at the time of the first flight.

The Block II engine also incorporated the advanced health management system on STS-117 in 2007. This on-board system could detect and mitigate anomalous high-pressure turbopump vibration behavior, and

the system further improved engine ascent safety by an additional 23%.

Summary

Another major SSME milestone took place in 2004 when the main engine passed 1,000,000 seconds in test and operating time. This unprecedented level of engine maturity over the preceding 3 decades established the main engine as one of the world's most reliable rocket engines, with a 100% flight safety record and a demonstrated reliability exceeding 0.9996 in over 1,000,000 seconds of hot-fire experience.



Chemochromic Hydrogen Leak Detectors

The Chemochromic Point Detector for sensing hydrogen gas leakage is useful in any application in which it is important to know the presence and location of a hydrogen gas leak.

This technology uses a chemochromic pigment and polymer that can be molded or spun into a rigid or pliable shape useable in variable-temperature environments including atmospheres of inert gas, hydrogen gas, or mixtures of gases. A change in the color of detector material reveals the location of a leak. Benefits of this technology include: temperature stability, from -75°C to 100°C (-103°F to 212°F); use in cryogenic applications; ease of application and removal; lack of a power requirement; quick response time; visual or electronic leak detection; nonhazardous qualities, thus environmentally friendly; remote monitoring capability; and a long shelf life. This technology is also durable and inexpensive.

The detector can be fabricated into two types of sensors—reversible and irreversible. Both versions immediately notify the operator of the presence of low levels of hydrogen; however, the reversible version does not require replacement after exposure. Both versions were incorporated into numerous polymeric materials for specific applications including: extruded tapes for wrapping around valves and joints suspected of leaking; injection-molded parts for seals, O-rings, pipe fittings, or plastic piping material; melt-spun fibers for clothing applications; and paint for direct application to ground support equipment. The versatility of the sensor for several different applications provides the operator with a specific-use safety notification while working under hazardous operations.



Hydrogen-sensing tape applied to the Orbiter midbody umbilical unit during fuel cell loading for STS-118 through STS-123 at Kennedy Space Center, Florida.

Hydrogen-sensing tape application at liquid hydrogen cross-country vent line flanges on the pad slope.



The First Human-Rated Reusable Solid Rocket Motor

The Space Shuttle reusable solid rocket motors were the largest solid rockets ever used, the first reusable solid rockets, and the only solids ever certified for crewed spaceflight. The closest solid-fueled rival—the Titan IV Solid Rocket Motor Upgrade—was known for boosting heavy payloads for the US Air Force and National Reconnaissance Organization. The motors were additionally known for launching the 5,586-kg (12,220-pound) Cassini mission on its 7-year voyage to Saturn. By contrast, the Titan booster was 76 cm (30 in.) smaller in diameter and 4.2 m (14 ft) shorter in length, and held only two-thirds of the amount of propellant.

In a class of its own, the Reusable Solid Rocket Motor Program was characterized from its inception by four distinguishing traits: hardware reusability, postflight recovery and analysis, a robust ground-test program, and a culture of continual improvement via process control.

The challenge NASA faced in developing the first human-rated solid rocket motor was to engineer a pair of solid-fueled rocket motors capable of meeting the rigorous reliability requirements associated with human spaceflight. The rocket motors would have to be powerful enough to boost the shuttle system into orbit. The motors would also need to be robust enough to meet stringent reliability requirements and survive the additional rigors of re-entry into Earth's atmosphere and subsequent splashdown, all while being reusable. The prime contractor—Morton Thiokol, Utah—completed its



The two shuttle reusable solid rocket motors, which stood more than 38 m (126 ft) tall, harnessed 29.4 meganewtons (6.6 million pounds) of thrust. The twin solid-fueled rockets provided 80% of the thrust needed to achieve liftoff.

first full-scale demonstration test within 3 years.

NASA learned a poignant lesson in the value of spent booster recovery and inspection with the Challenger tragedy in January 1986. The postflight condition of the hardware provided valuable information on the health of the design and triggered a redesign effort that surpassed, in magnitude and complexity, the original development program.

For the substantial redesign that occurred between 1986 and 1988, engineers incorporated lessons learned from the first 25 shuttle flight booster sets. More than 100 tests, including five full-scale ground tests, were conducted to demonstrate the strength of the new design. Flaws were deliberately manufactured into the final test motor to check redundant systems.

The redesigned motors flew for the first time in September 1988 and performed flawlessly.

A Proven Design

To construct the reusable solid rocket motor, four cylindrical steel segments—insulated and loaded with a high-performance solid propellant—were joined together to form what was essentially a huge pressure vessel and combustion chamber. The segmented design provided maximum flexibility in motor fabrication, transportation, and handling. Each segment measured 3.7 m (12 ft) in diameter and was forged from D6AC steel measuring approximately 1.27 cm (0.5 in.) in thickness.

Case integrity and strength were maintained during flight by insulating the case interior. The insulating liner was a fiber-filled elastomeric (rubber-like) material applied to the interior of the steel cylinders. A carefully formulated tacky rubber bonding layer—or “liner”—was applied to the rubber insulator surface to facilitate a strong bond with the propellant.

Producing an accurate insulating layer was critical. Too little insulation, and

the steel could be heated and melted by the 2,760°C (5,000°F) combustion gases. Too much insulation, and weight requirements were exceeded. Engineers employed sophisticated design analysis and testing to optimize this balance between protection and weight. By design, much of the insulation was burned away during the 2 minutes of motor operation.


The propellant was formulated from three major ingredients: aluminum powder (fuel); ammonium perchlorate (oxidizer); and a synthetic polymer binding agent. The ingredients were batched, fed into large 2,600-L (600-gal) mix bowls, mixed, and tested before being poured into the insulated and lined segments. Forty batches were produced to fill each case segment. The propellant mixture had an initial consistency similar to that of peanut butter, but was cured to a texture and color that resembled a rubber pencil eraser—strong, yet pliable. The propellant configuration or “shape” inside each segment was carefully designed and cast to yield the precise thrust trace upon ignition.

Once each segment was insulated and cast with propellant and finalized, the segments were shipped from ATK’s manufacturing facility in Utah to Kennedy Space Center (KSC) in Florida, on specially designed, heavy-duty covered rail cars. At KSC, they were stacked and assembled into the flight configuration.

The segments were joined together with tang/clevis joints pinned in 177 locations and sealed with redundant O-rings. Each joint, with its redundant seals and multiple redundant seal protection features, was pressure checked during assembly to ensure a good pressure seal.



Reusable Solid Rocket Motor Propellant Configuration



The four primary propulsion segments that comprised the reusable solid rocket motor were manufactured individually then assembled for launch. Each segment was reusable and designed for a service life of up to 20 flights.

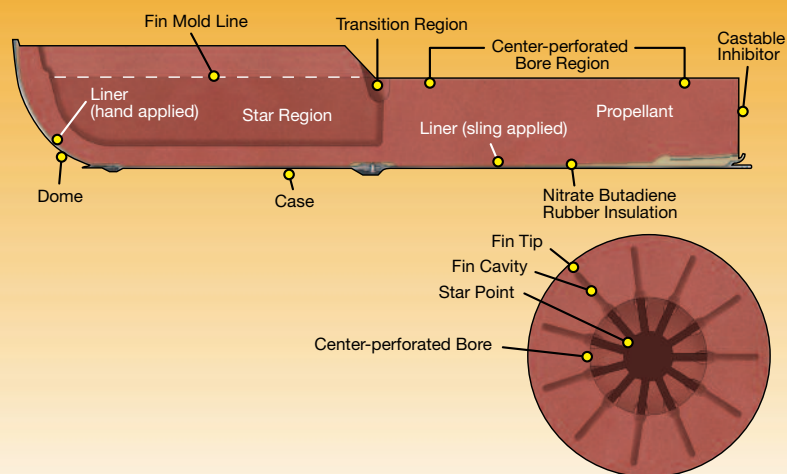
© ATK. All rights reserved.

An igniter was installed in the forward end of the forward segment—at the top of the rocket. The igniter was essentially a smaller rocket motor that fired into the solid rocket motor to ignite the main propellant grain. Design and manufacture closely mirrored the four main segments.

The nozzle was installed at the aft end of the aft segment, at the bottom of the rocket. The nozzle was the “working” component of the rocket in which hot exhaust gases were accelerated and directed to achieve performance requirements and vehicle control.

The nozzle structure consisted of metal housings over which were bonded layers of carbon/phenolic and silica/phenolic materials that protected the metal structure from the searing exhaust gases by partially decomposing and ablating. A flexible bearing, formed with vulcanized rubber and steel, allowed for nozzle maneuverability up to 8 degrees in any direction to steer the shuttle during the first minutes of flight.

Forward Segment Propellant Grain Configuration



The forward propulsion segment featured a unique grain pattern designed to yield the greatest thrust when it was needed most—on ignition.

© ATK. All rights reserved.

Engineers employed significant analysis and testing to develop a reliable and efficient nozzle capable of being manufactured. The nozzle flexible bearing—measuring up to 2.35 m (92.4 in.) at its outside diameter—was an example of one component that required multiple processing iterations to ensure the manufactured product aligned with design requirements.

NASA enhanced the nozzle design following the Challenger accident when severe erosion on one section of the nozzle on one motor was noted through postflight analysis. While the phenolic liners were designed to erode smoothly and predictably, engineers found—at certain ply orientations—that internal stresses resulting from exposure to hot

gases exceeded the material strength. Under such stress, the hot charred material had the potential to erode erratically and jeopardize component integrity. Engineers modified nozzle ply angels to reduce material stress, and this condition was successfully eliminated on all subsequent flights.



Technicians shown installing igniter used to initiate the propellant burn in a forward motor segment. The igniter was a small rocket motor loaded with propellant that propagated flame down the bore of the motor.

The Reusable Rocket

All metal hardware—including structures from the case, igniter, safe-and-arm device, and nozzle—were designed to support up to 20 shuttle missions. This was unique to the reusable solid rocket motor. Besides the benefits of conservation and affordability, the ability to recover the motors allowed NASA to understand exactly how the components performed in flight. This performance analysis provided a wealth of valuable information and created a synergy to drive improvements in motor performance, implemented through motor manufacturing and processing.

This recovery and postflight capability was particularly important for the long-term Space Shuttle Program since,

over time, changes were inevitable. Change to design or process became mandatory as a result of factors such as material/vendor obsolescence or new environmental regulations.

Changing Processes

During a 10-year period beginning in the mid 1990s, for example, more than 100 supplier materials used to produce the reusable solid rocket motor became obsolete. The largest contributing factor stemmed from supplier economics, captured in three main scenarios. First, suppliers changed their materials or processes. Second, suppliers consolidated operations and either discontinued or otherwise modified their materials. Third, the materials were simply no longer available from subtier vendors.

US environmental regulations, such as the requirement to phase out the use of ozone-depleting chemicals, were an additional factor. Methyl chloroform, for example, was a solvent used extensively in hardware processing. A multimillion-dollar effort was launched within NASA and ATK to eventually eliminate methyl chloroform use altogether in motor processing. Eight alternate materials were selected following thorough testing and analysis to ensure program performance was not compromised.

New Technology

Advancements in technology that occurred during the decades-long program were a further source of change. Engineers incorporated new technologies into motor design and processing as the technology could be proven. Incorporating braided carbon fiber material as a thermal barrier in the nozzle-to-case joint is one example.

Postflight Analysis

The ability to closely monitor flight performance through hands-on postflight analysis—after myriad material, design, and process changes—was only possible by virtue of the motor's reusable nature.


Developing methods to scrutinize and recertify spent rocket motor hardware that had raced through the stratosphere at supersonic speeds was new. NASA had the additional burden of working with components that had experienced splashdown loads and were subsequently soaked in corrosive saltwater prior to retrieval.

In the early days of the program, NASA made significant efforts in identifying relevant evaluation criteria and establishing hardware assessment methods. A failure to detect hardware stresses and material weaknesses could result in an unforgivable catastrophic event later on. The criteria used to evaluate the first motors and the accompanying data collected would also become the benchmark from which future flights would be measured. Included in the evaluation criteria were signs of case damage or material loss caused by external debris; integrity of major components such as case segments, nozzle and igniter; and fidelity of insulation, seals, and joints.

Inspection and documentation of retrieved hardware occurred in two parts of the country: Florida, where the hardware was retrieved; and Utah, where it underwent in-depth inspection and refurbishment. On recovery, a team of 15 motor engineers conducted what was termed an "open assessment," primarily focusing on exterior components. After retrieval, teams of specialists rigorously dissected, measured, sampled, and assessed joints,



Field Joint Comparison for Use on Reusable Solid Rocket Motor



Reusable solid rocket motors incorporated significant improvements over the earlier shuttle motors in the design of the joints between the main segments. Redesign of this key feature was part of the intensive engineering redesign and demonstration feat accomplished following the Challenger accident. The result was a fail-safe joint/seal configuration that, with continued refinement, had a high demonstrated reliability. Each joint, with its redundant seals and multiple redundant seal protection features, could be pressure checked during assembly to ensure a good pressure seal was achieved. A similar design approach was implemented on the igniter joints during that same time period.

© ATK. All rights reserved.

bondlines, ablatives, fasteners, and virtually all remaining flight hardware. Engineers promptly evaluated any significant observations that could affect the orbiting vehicle or the next motor launch sets.

Before the motor was returned to the flight inventory, the recovered metal parts were inspected for corrosion, deformations, cracks, and other potential damage. Dimensional measurement

data were fed into a system-wide database containing documentation dating back to the program's inception. The wealth of information available for performance trend analysis was unmatched by any other solid rocket motor manufacturing process in the world. Gates and checks within the system ensured the full investigation of any anomalies to pinpoint root cause and initiate corrective action.

The postflight analysis program collected the actual flight performance data—most of which would not have been available if the motors had not been recovered.

Through this tightly defined process, engineers were able to address the subtle effects that are often a result of an unintended drift in the manufacturing process or new manufacturing materials introduced into the process. The



process addressed these concerns in the incipient phase rather than allowing for a potentially serious issue to escalate undetected. The ultimate intangible benefit of this program was greater reliability, as demonstrated by the following two examples.

Postflight assessment of nozzle bondlines was a catalyst to augment adhesive bonding technology and substantially improve hardware quality and reliability. Storage controls for epoxy adhesives were established in-house and with adhesive suppliers. Surface preparation, cleanliness, adhesive primer, and process timelines were established. Adhesive bond quality and robustness were increased by an order of magnitude.

Postflight inspections also occasionally revealed gas paths through the nozzle-to-case joint polysulfide thermal barrier that led to hot gas impingement on the wiper O-ring—a structure protecting the primary O-ring from thermal damage. While this condition did not pose a flight risk, it did indicate performance failed to meet design intent. The root cause: a design that was impossible to manufacture perfectly every time. Engineers resolved this concern by implementing a nozzle-to-case joint J-leg design similar to that successfully used on case field joints and igniters.

Robust Systems Testing

The adage “test before you fly,” adopted by the Space Shuttle Program, was the standard for many reusable solid rocket motor processes and material, hardware, and design changes. What ATK, the manufacturer, was able to learn from the vast range of data collected and processed through preflight and ground testing ensured



In Utah, rigorous test program included 53 reusable solid rocket motor ground tests between 1977 and 2010. Spectators flocked by the thousands to witness firsthand the equivalent of 15 million horsepower safely unleashed from a vantage point of 2 to 3 km (1 to 2 miles) away.

the highest levels of dependability and safety for the hardware. Immediate challenges posed by the 570-metric-ton (1.2-million-pound) motor included handling, tooling, and developing a 17.8-meganewton (4,000,000-pound-force) thrust-capable ground test stand; and designing a 1,000-channel data handling system as well as new support systems, instrumentation capability, data acquisition, and countdown procedures.

Hot-fire testing of full-scale rocket motors in the Utah desert became a hallmark of the reusable solid rocket motor development and sustainment program. Individual motor rockets were fired horizontally, typically once or twice a year, lighting up the mountainside with the brightness of a blazing sun, even in broad daylight.

Following a test firing, quick-look data were available within hours. Full data analyses required several months.

On average, NASA collected between 400 and 700 channels of data for each test. Instrumentation varied according to test requirements but typically

included a suite of sensors not limited to accelerometers, pressure transducers, calorimeters, strain gauges, thermocouples, and microphones. Beyond overall system assessment and component qualification, benefits of full-scale testing included the opportunity to enhance engineering expertise and predictive skills, improve engineering techniques, and conduct precise margin testing. The ability to tightly measure margins for many motor process, material, components, and design parameters provided valuable verification data to demonstrate whether even the slightest modification was safe for flight.

Quick-look data revealed basic ballistics performance—pressure and thrust measurements—that could be compared with predicted performance and historic data for an initial assessment.

Full analysis included scrutiny of all data recorded during the actual test as well as additional data gathered from visual inspections and measurements of disassembled hardware, similar



to that of postflight inspection. Engineers assessed specific data tied to test objectives. When qualifying a new motor insulation, for example, posttest inspection would additionally include measurements of remaining insulation material to calculate the rate of material loss.

Subscale propellant batch ballistics tests, environmental conditioning testing, vibration tests, and custom sensor development and data acquisition were also successful components of the program to provide specific reliability data.

Culture of Continual Improvement

The drive to achieve 100% mission success, paired with the innovations of pre- and postflight testing that allowed performance to be precisely quantified, resulted in an operating culture in which the bar was continually raised.

Design and processing improvements were identified, pursued, and implemented through the end of the program to incrementally reduce risk and waste. Examples of relatively late program innovations included: permeable carbon fiber rope as a thermal protection element in various nozzle and nozzle/case joints; structurally optimized bolted joints; reduced stress forward-grain fin transition configuration; and improved adhesive bonding systems.

This culture, firmly rooted in the wake of the Challenger accident, led to a comprehensive process control program with systems and tools to ensure processes were appropriately defined, correctly performed, and adequately maintained to guarantee reliable and repeatable product performance.

Noteworthy elements of the motor process control program included an extensive chemical fingerprinting program to analyze and monitor the quality of vendor-supplied materials, the use of statistical process control to better monitor conditions, and the comprehensive use of witness panels—product samples captured from the live manufacturing process and analyzed to validate product quality.

With scrupulous process control, ATK and NASA achieved an even greater level of understanding of the materials and processes involved with reusable solid rocket motor processing. As a result, product output became more consistent over the life of the program. Additionally, partnerships with vendors and suppliers were strengthened as increased performance measurement and data sharing created a win-win situation.

An Enduring Legacy

The reusable solid rocket motor was more than an exceptional rocket that safely carried astronauts and hundreds of metric tons of hardware into orbit for more than 25 years. Throughout the Reusable Solid Rocket Motor Program, engineers and scientists generated the technical know-how in design, test, analysis, production, and process control that is essential to continued space exploration. The legacy of the first human-rated reusable solid rocket motor will carry on in future decades. In the pages of history, the shuttle reusable solid rocket motor will be known as more than a stepping-stone. It will also be regarded as a benchmark by which future solid-propulsion systems will be measured.

Orbital Propulsion Systems— Unique Development Challenges

Until the development of the Space Shuttle, all space vehicle propulsion systems were expendable. Influenced by advances in technologies and materials, NASA decided to develop a reusable propulsion system. Although reusability saved overall costs, maintenance and turnaround costs offset some of those benefits.


NASA established a general redundancy requirement of fail operational/fail safe for these critical systems: Orbital Maneuvering System, Reaction Control System, and Auxiliary Power Unit. In addition, engineers designed the propulsion systems for a life of 100 missions or 10 years combined storage and operations. Limited refurbishment was permitted at the expense of higher operational costs.

Orbital Maneuvering System

The Orbital Maneuvering System provided propulsion for the Orbiter during orbit insertion, orbit circularization, orbit transfer, rendezvous, and deorbit. NASA faced a major challenge in selecting the propellant. The agency originally chose liquid oxygen and liquid hydrogen propellants. However, internal volume constraints could not be met for a vehicle configuration that provided a payload of 22,680 kg (50,000 pounds) in a bay measuring 4.6 m (15 ft) in diameter and 18.3 m (60 ft) in length. This, coupled with concerns regarding complexity of cryogenic propellants, led to the consideration of storable hypergolic propellants.



Orbital Maneuvering System/Reaction Control System



NASA ultimately selected monomethylhydrazine as the fuel and nitrogen tetroxide as the oxidizer for this system. As these propellants were hypergolic—they ignited when coming into contact with each other—no ignition device was needed. Both propellants remained liquid at the temperatures normally experienced during a mission. Electrical heaters prevented freezing during long periods in orbit when the system was not in use.

Modular Design Presents Obstacles for Ground Support

Trade studies and design approach investigations identified challenges and solutions. For instance, cost and weight could be reduced with a common integrated structure for the Orbital Maneuvering System and Reaction Control System. This integrated structure was combined with the selection of nitrogen tetroxide and monomethylhydrazine propellants.

Thus, NASA adopted an interconnect system in which the Reaction Control System used Orbital Maneuvering System propellants because of cost, weight, and lower development risk.

Disadvantages of a storable propellant system were higher maintenance requirements resulting from their corrosive nature and hazards to personnel exposed to the toxic propellants. NASA partially addressed these considerations by incorporating the Orbital Maneuvering System into a removable modular pod. This allowed maintenance and refurbishment of those components exposed to hypergols to be separated from other turnaround activities.

For ground operations, it was not practical to remove modules for each turnaround activity, and sophisticated equipment and processes were required for servicing between flights. Fluid and gas connections to the propellants and pressurants used quick disconnects to allow servicing on the launch pad, in Orbiter processing facilities, and in the hypergolic maintenance facility. However, quick disconnects occasionally caused problems, including leakage that damaged Orbiter thermal tiles.

Engineers tested and evaluated many ground support equipment design concepts at the White Sands Test Facility (WSTF). In particular, they tested, designed, and built the equipment used to test and evaluate the propellant acquisition screens inside the propellant tanks before shipment to Kennedy Space Center for use on flight vehicles. The Orbital Maneuvering System/Reaction Control System Fleet Leader Program used existing qualification test articles to detect and evaluate “life-dependent” problems before these problems affected the



shuttle fleet. This program provided a test bed for developing and evaluating ground support equipment design changes and improving processes and procedures. An example of this was the Reaction Control System Thruster Purge System, which used low-pressure nitrogen to prevent propellant vapors from accumulating in the thruster chamber. This WSTF-developed ground support system proved beneficial in reducing the number of in-flight thruster failures.

Additional Challenges

Stable combustion was a concern for NASA. In fact, stable combustion has always been the most expensive schedule-constraining development issue in rocket development. For the Orbital Maneuvering System engine, engineers investigated injector pattern designs combined with acoustic cavity concepts. In propulsion applications with requirements for long-duration firings and reusability, cavities had an advantage because they were easy to

cool and therefore less subject to failure from either burnout or thermal cycling.

To accomplish precise injector fabrication, engineers implemented platelet configuration. The fuel and oxidizer flowed through the injector and impinged on each other, causing mixing and combustion. Platelet technology, consisting of a series of thin plates manufactured by photo etching and diffusion bonded together,



Henry Pohl
*Director of Engineering at Johnson Space Center
(1986-1993).*

"To begin to understand the challenges of operating without gravity, imagine removing the commode from your bathroom floor, bolting it to the ceiling. And then try to use it. You would then have a measure of the challenges facing NASA."

eliminated mechanical manufacturing errors and increased injector life and combustion efficiency.

The combustion chamber was regenerative-cooled by fuel flowing in a single pass through non-tubular coolant channels. The chamber was composed of a stainless-steel liner, an electroformed nickel shell, and an aft flange and fuel inlet manifold assembly. Its structural design was based on life

Formation of Metal Nitrates Caused Valve Leaks

Being the first reusable spacecraft—and in particular, the first to use hypergolic propellants—the shuttle presented technical challenges, including leaky and sticky propellant valves in the Reaction Control System thrusters. Early in the program, failures in this system were either an oxidizer valve leak or failure to reach full chamber pressure within an acceptable amount of time after the thruster was commanded on. NASA attributed both problems to the buildup of metal nitrates on and around the valve-sealing surfaces.

Metal nitrates were products of iron dissolved in the oxidizer when purchased and iron and nickel that were leached out of the ground and flight fluid systems. When the oxidizer was exposed to reduced pressure or allowed to evaporate, metal nitrates precipitated out of solution and contaminated the valve seat.


Subsequent valve cycling caused damage to the Teflon® valve seat, further exacerbating the leakage until sufficient nitrate deposition resulted in "gumming" up the valve. At that point, the valve was either slow to operate or failed to operate.

Multiple changes reduced the metal nitrate problem but may have contributed to fuel valve seat extrusion, which manifested years later. The fuel valve extrusion was largely attributed to the use of throat plugs. These plugs trapped oxidizer vapor leakage in the combustion chamber, which subsequently reacted at a low level of fuel that had permeated the Teflon® fuel valve seat. This problem was successfully addressed with the implementation of the NASA-developed thruster nitrogen purge system, which kept the thruster combustion chamber relatively free of propellant vapors.



An Ordinary Solution to the Extraordinary Challenge of Rain Protection

During operations, Orbiter engines needed rain protection after the protective structure was moved away and protective ground covers were removed. This requirement protected the three upward-facing engines and eight of the left-side engines from rainwater accumulation on the launch pad. The up-firing engine covers had to prevent water accumulation that could freeze in the injector passages during ascent. The side-firing engine covers prevented water from accumulating in the bottom of the chamber and protected the chamber pressure sensing ports. Freezing of accumulated water during ascent could block the sensing port and cause the engine to be declared “failed off” when first used. The original design concept allowed for Teflon® plugs installed in the engine throats and a combination of Teflon® plugs tied to a Teflon® plate that covered the nozzle exit. This concept added vehicle weight, required special procedures to eject the plugs in flight, and risked accidental ejection in ascent that could damage tiles. The solution used ordinary plastic-coated freezer paper cut to fit the exit plane of the nozzle. Tests proved this concept could provide a reliable seal under all expected rain and wind conditions. The covers were low cost, simple, and added no significant weight. The thruster rain cover material was changed to Tyvek® when NASA discovered pieces of liberated plastic-coated paper beneath the cockpit window pressure seals. The new Tyvek® covers were designed to release at relatively low vehicle velocity so that the liberated covers did not cause impact damage to windows, tile, or any other Orbiter surface.



Tyvek® covers shown installed on forward Reaction Control System thrusters (top) and a typical cover (right). Note that the covers were designed to fit certain thruster exit plane configurations.



cycle requirements, mechanical loads, thrust and aerodynamic loading on the nozzle, ease of fabrication, and weight requirements.

The nozzle extension was radiation cooled and constructed of columbium metal consistent with experience gained during the Apollo Program. The mounting flange consisted of a bolt ring, made from a forging and a tapered section, that could either be spun or made from a forging. The forward and aft sections were made from two panels each. This assembly was bulge formed to the final configuration and the stiffening rings were attached by welding. The oxidation barrier diffusion operation was done after machining was completed.

A basic design challenge for the bipropellant valve was the modular valve. The primary aspect of the assembly design was modularization, which reduced fabrication problems and development time and allowed servicing and maintenance goals to be met with lower inventory.


NASA Seeks Options as Costs Increase

The most significant lesson learned during Orbital Maneuvering System development was the advantage of developing critical technologies before initiating full-scale hardware designs. The successful completion of predevelopment studies not only reduced total costs, also it minimized schedule delays.

In the 1980s, NASA began looking for ways to decrease the cost of component refurbishment and repair. NASA consolidated engineering, evaluation, and repair capabilities for many components, and reduced overall costs. Technicians serviced, acceptance



Forward Reaction Control System



Forward Reaction Control System on Discovery.

tested, and prepared all hypergolic wetted components for reinstallation on the vehicles.

Reaction Control System

The Reaction Control System provided propulsive forces to control the motion of the Orbiter for attitude control, rotational maneuvers, and small velocity changes along the Orbiter axes. The

requirement of a fail-operational/fail-safe design introduced complexity of additional hardware and a complex critical redundancy management system. The reuse requirement posed problems in material selection and compatibility, ground handling and turnaround procedures, and classical wear-out problems. The requirement for both on-orbit operations and re-entry into Earth's atmosphere complicated

propellant tank acquisition system design because of changes in the gravitational environment.

NASA Makes Effective Selections

As with the Orbital Maneuvering System, propellant selection was important for the Reaction Control System. NASA chose a bipropellant of monomethylhydrazine and nitrogen

Low Temperatures, Increased Leakage, and a Calculated Solution

Some primary thruster valves could leak when subjected to low temperature. NASA discovered this problem when they observed liquid dripping from the system level engines during a cold environment test. The leakage became progressively worse with increased cycling. Continued investigation

indicated that tetrafluoroethylene Teflon® underwent a marked change in the thermal expansion rate in a designated temperature range. Because machining, done as a part of seat fabrication, was accomplished in this temperature range, some parts had insufficient seat material exposed at reduced temperatures.


To reduce susceptibility to cold leakage, engineers machined Teflon® at 0°C (32°F) to ensure uniform dimensions with adequate seat material exposed at reduced temperatures and raised the thruster heater set points to maintain valve temperature above 16°C (60°F).

Cracks Prompt Ultrasonic Inspection

Late in the Space Shuttle Program, NASA discovered cracks in a thruster injector. The thruster was being refurbished at White Sands Test Facility (WSTF) during the post-Columbia accident Return to Flight time period. The cracks were markedly similar to those that had occurred in injectors in 1979 and again in 1982.

These earlier cracks were discovered during manufacturing of the thrusters and occurred during the nozzle insulation bake-out process. Results from the laboratory testing indicated that cracks were developed due to chemical processing and manufacturing. In addition to using leak testing to screen for injector cracking, NASA engineers developed and implemented an ultrasonic inspection procedure to screen for cracks that measured less than the injector wall thickness.

The marked similarity of the crack location and crack surface appearance strongly suggested the WSTF-discovered cracks were due to the original equipment manufacturing process and were not flight induced or propagated. Laboratory tests and analyses confirmed that those cracks were induced in manufacturing. The cracks had not grown significantly over the years of the thruster's use and its many engine firings. Laboratory nondestructive testing showed that the original ultrasonic inspection process was not very reliable and it was possible that manufacturing-induced cracks could



Reaction Control System thruster cross sections showing the crack location and its actual surface appearance.

escape detection and cracked thrusters could have been placed in service. The fact that there was no evidence of crack growth associated with the WSTF-discovered

cracks due to the service environment was a significant factor in the development of flight rationale for the thrusters.



tetroxide system, which allowed for integration of this system with the Orbital Maneuvering System. This propellant combination offered a favorable weight tradeoff, reasonable development cost, and minimal development risk.

NASA selected a screen tank as a reusable propellant supply system to provide gas-free propellants to the thrusters. Screen tanks worked by using the surface tension of the liquid to form a barrier to the pressurant gas. The propellant acquisition device was made of channels covered with a finely woven steel mesh screen. Contact with liquid wetted the screen and surface tension of the liquid prevented the passage of gas. The strength of the liquid barrier was finite. The pressure differential at which gas would be forced through the wetted screen was called the “bubble point.” When the bubble point was exceeded, the screen broke down and gas was transferred. If the pressure differential was less than the bubble point, gas could not penetrate the liquid barrier and only liquid was pulled through the channels. NASA achieved their goal in designing the tank to minimize the pressure loss while maximizing the amount of propellant expelled.

Several Reaction Control System component failures were related to nitrate contamination. Storage of oxidizer in tanks and plumbing that contained iron caused contamination in the propellant. This contamination formed a nitrate that could cause valve leakage, filter blockage, and interference in sliding fits. The most prominent incident was the failure of a ground half-quick disconnect to close, resulting in an oxidizer spill on the launch pad. NASA implemented

a program to determine the parameters that caused the iron nitrate formation and implement procedures to prevent its formation in the future. This resulted in understanding the relationship between iron, water, nitric oxide content, and nitrate formation. The agency developed production and storage controls as well as filtration techniques to remove the iron, which resolved the iron nitrate problem.

Auxiliary Power Unit

The Auxiliary Power Unit generated power to drive hydraulic pumps that produced pressure for actuators to control the main engines, aero surfaces, landing gear, brakes, and nose wheel steering. The Auxiliary Power Unit shared common hardware and systems with the Hydraulic Power Unit used on the solid rocket motors. The shuttle needed a hydraulic power unit that could operate from zero to three times gravity, at vacuum and sea-level pressures, from -54°C to 107°C (-65°F to 225°F), and be capable of restarting. NASA took the basic approach of using a small, high-speed, monopropellant-fuel, turbine-powered unit to drive a conventional aircraft-type hydraulic pump.

If the Auxiliary Power Unit was restarted before the injector cooled to less than 204°C to 232°C (400°F to 450°F), the fuel would thermally decompose behind the injector panels and damage the injector and the Gas Generator Valve Module. Limited hot-restart capability was achieved by adding an active water cooling system to the gas generator to be used only for hot restarts. This system injected water into a cavity within the injector. The

steam generated was vented overboard. Use of this system enabled restarts at any time after the cooling process, which required a 210-second delay.

Improved Machining and Manufacturing Solves Valve Issue

Development of a reliable valve to control fuel flow into the gas generator proved to be one of the most daunting tasks of the propulsion systems. The valve was required to pulse fuel into the gas generator at frequencies of 1 to 3 hertz. Problems with the valve centered on leakage and limited life due to wear and breakage of the tungsten carbide seat. NASA's considerable effort in redesigning the seat and developing manufacturing processes resulted in an intricate seat design with concentric dual sealing surfaces and redesigned internal flow passages. The seat was diamond-slurry honed as part of the manufacturing process to remove the recast layer left by the electro-discharge machining. This recast layer was a source of stress risers and was considered one of the primary factors causing seat failure. The improved design and machining and manufacturing processes were successful.

Additional Challenges and Subsequent Solutions

During development testing of the gear box, engineers determined that the oil pump may not function satisfactorily on orbit due to low pressure. It became necessary to provide a fluid for the pump to displace to assure the presence of oil at the inlet and to have a mechanism to provide needed minimum pressure at startup and during operation.

The Auxiliary Power Unit was designed with a turbine wheel radial containment ring and a blade tip seal and rub ring to safely control failures of the high-speed assembly. The containment ring was intended to keep any wheel fragments from leaving the Auxiliary Power Unit envelope. NASA

provided safety features that would allow operation within the existing degree of containment. The agency used an over-speed safety circuit to automatically shut down a unit at 93,000 revolutions per minute. To provide further insurance against wheel failure, NASA imposed stringent flaw detection


inspections. With these controls, results of fracture mechanics analyses showed the theoretical life to be 10 times the 100-mission requirement.

With these improvements, the Auxiliary Power Unit demonstrated success of design and exhibited proven durability, performance, and reusability.

NASA Encounters Obstacle Course in Turbine Wheel Design

The space agency faced multiple challenges with the development of the turbine wheel. Aerodynamically induced high-cycle fatigue caused cracking. Analysis indicated this part of the blade could be removed with a small chamfer at the blade tip without significant effect on performance. This cracking problem was resolved by careful design and control of electromechanical machining.

The shroud cracking problem was related to material selection and the welding process. Increased strength and weld characteristics were achieved by changing the shroud material. Engineers developed a controlled electron beam weld procedure to ensure no overheating of the shroud. These actions eliminated the cracking problem.






Stress Corrosion and Propellant Ignition

One of the most significant Auxiliary Power Unit problems occurred during the STS-9 (1983) mission when two of the three units caught fire and detonated. Postflight analysis indicated the presence of hydrazine leaks in Auxiliary Power Units 1 and 2 when they were started for re-entry while still in orbit. The leaking hydrazine subsequently ignited and the resulting fire

overheated the units, causing the residual hydrazine to detonate after landing. The fire investigation determined the source of the leaks to be nearly identical cracks in the gas generator injector tubes in both units. Laboratory tests further determined that the injector tube cracks were due to stress corrosion from ammonium hydroxide vapors generated by decomposition of

hydrazine in the catalyst bed after Auxiliary Power Unit shutdown.

Initial corrective actions included removal of the electrical machined recast layer on the tube inside diameter and an improved assembly of the injector tube. Later, resistance to stress corrosion and general corrosion was further improved by chromizing the injector tubes.



Summary

The evolution of orbital propulsion systems for the Space Shuttle Program began with Apollo Program concepts, expanded with new

technologies required to meet changing requirements, and continued with improvements based on flight experience. The design requirements for 100 missions, 10 years, and reuse presented challenges not previously

encountered. In addition, several problems were not anticipated. NASA met these challenges, as demonstrated by the success of these systems.



Pioneering Inspection Tool

Contamination Scanning of Bond Surfaces


Bonding thermal insulation to metal case surfaces was a critical process in solid rocket motor manufacturing during the Space Shuttle Program. Surfaces had to be immaculately clean for proper adherence. The steel alloy was susceptible to corrosion and was coated with grease for protection during storage. That grease, and the solvents to remove it, became potential contaminants.

The improvement of contamination inspection techniques was initiated in the late 1980s. The development of a quantitative and recordable inspection technique was based on the physics of optically stimulated electron emission (photoelectric effect) technology being developed at NASA's Marshall Space Flight Center at the time.

Fundamentally, incident ultraviolet light excites and frees electrons from the metal surface. The freed electrons having a negative charge are attracted to a positively charged collector ring in the "Con Scan" (short for Contamination Scanning) sensor. When contamination exists on a metal surface, the amount of ultraviolet radiation that reaches the surface is reduced. In turn, the current is reduced, confirming the presence of a contaminant.

Approximately 90% of each reusable solid rocket motor barrel assembly was inspected using automated Con Scan before bond operations. Technicians mounted the sensor on a robotic arm, which allowed longitudinal translation of the sensor as the barrel assembly rotated on a turntable. Inspection results were mapped, showing color-coded contamination levels (measured current) vs. axial and circumferential locations on the case inner diameter. Color coding made acceptable and rejected areas visually apparent.

By pioneering optically stimulated electron emission technology, which was engineered into a baseline inspection tool, the Space Shuttle Program significantly improved contamination control methods for critical bonding applications.



Inspection technology capitalizing on the photoelectric effect provided significant benefits over the traditional method of visual inspection using handheld black lights. The technology was developed through a NASA/industry partnership managed by Marshall Space Flight Center. Specific benefits included increased accuracy in contamination detection and an electronic data record for each hardware inspection.

Propulsion Systems and Hazardous Gas Detection

Shuttle propulsion had hazardous gases requiring development of detection systems including purged compartments. This development was based on lessons learned from the system first used during Saturn I launches.

NASA performed an exhaustive review of all available online monitoring mass spectrometry technology for the shuttle. The system the agency selected for the prototype Hazardous Gas Detection System had an automated high-vacuum system, a built-in computer control interface, and the ability to meet all program-anticipated detection limit requirements.

The instrument arrived at Kennedy Space Center (KSC) in December 1975 and was integrated into the sample delivery subsystem, the control and data subsystem, and the remote control subsystem designed by KSC. Engineers extensively tested the unit for functionality, detection limits and dynamic range, long-term drift, and other typical instrumental performance characteristics. In May 1977, KSC shipped the prototype Hazardous Gas Detection System to Stennis Space Center to support the shuttle main propulsion test article engine test firings. The system remained in use at Stennis Space Center for 12 years and supported the testing of upgraded engines.

The first operational Hazardous Gas Detection System was installed for the system on the Mobile Launch



Platform-1 during the late summer of 1979. Checkout and operations procedure development and activation required almost 1 year, but the system was ready to support initial purge activation and propellant loading tests in late 1980. A special test in which engineers introduced simulated leaks of hydrogen and oxygen into the Orbiter payload bay, lower midbody, aft fuselage, and the External Tank intertank area represented a significant milestone. The system accurately detected and measured gas leaks.

After the new system's activation issues were worked out, it could detect and measure small leaks from the Main Propulsion System. The Hazardous Gas Detection System did not become visible until Space Transportation System (STS)-6—the first launch of the new Orbiter Challenger—during a flight readiness test. In this test, the countdown would proceed normally to launch time, the Orbiter main engines would ignite, but the Solid Rocket Booster engines would not ignite and the shuttle would remain bolted to the launch pad during a 20-second firing of the main engines. The STS-1 firing test for Columbia had proceeded normally, but during Challenger's firing test, the Hazardous Gas Detection System detected a leak exceeding 4,000 parts per million. Rerunning the firing test and performing further leak hunting and analysis revealed a number of faults in the main engines. The manager for shuttle operation propulsion stated that all the money spent on the Hazardous Gas Detection System, and all that would ever be spent, was paid for in those 20 seconds when the leak was detected.

Originally, NASA declined to provide redundancy for the Hazardous Gas Detection System due to a lack of a launch-on-time requirement; however, the agency subsequently decided that redundancy was required. After a detailed engineering analysis followed by lab testing of candidate mass spectrometers, the space agency selected the PerkinElmer MGA-1200 as the basis of the backup Hazardous Gas Detection System. This backup was an ion-pumped, magnetic-sector, multiple-collector mass spectrometer widely used in operating rooms and industrial plants. Although the first systems were delivered in late 1985, full installation on all mobile launch platforms did not occur until NASA completed the Return to Flight activities following the Challenger accident in 1986.

In May 1990, the Hazardous Gas Detection System gained attention once again when NASA detected a hydrogen leak in the Orbiter aft fuselage on STS-35. The space agency also detected a hydrogen leak at the External Tank to Orbiter hydrogen umbilical disconnect and thought that the aft fuselage leakage indication was due to hydrogen from the external leak migrating inside the Orbiter. Workers rolled STS-35 back into the Vertical Assembly Building and replaced the umbilical disconnect. Meanwhile, STS-38 had been rolled to the pad and leakage was again detected at the umbilical disconnect, but not in the aft fuselage. STS-38 was also rolled back, and its umbilical disconnect was replaced. The ensuing investigation revealed that manufacturing defects in both units caused the leaks, but not before STS-35 was back on the pad.

During launch countdown, NASA detected the aft fuselage hydrogen leak. It was then apparent that STS-35 had experienced two separate leaks. The Space Shuttle Program director appointed a special tiger team to investigate the leak problem. This team suspected that the Hazardous Gas Detection System was giving erroneous data, and brought 10 experts from Marshall Space Flight Center to assess the system design. KSC design engineering provided an in-depth, 2-week description of the design and performance details of both the Hazardous Gas Detection System and the backup system. The most compelling evidence of the validity of the readings was that both systems, which used different technology, had measured identical data, and both systems had recorded accurate calibration data before and after leakage detection. After a series of mini-tanking tests—each with increased temporary instrumentation—engineers located and repaired the leak, and STS-35 lifted off for a successful mission on December 9, 1990.

The Hazardous Gas Detection System and backup Hazardous Gas Detection System continued to serve the shuttle until 2001, when both systems were replaced with Hazardous Gas Detection System 2000—a modern state-of-the-art system with a common sampling system and identical twin quadrupole mass spectrometers from Stanford Research Institute. The Hazardous Gas Detection System served for 22 years and the backup Hazardous Gas Detection System served for 15 years.



Thermal Protection Systems

Introduction

Gail Chapline

Orbiter Thermal Protection System

Alvaro Rodriguez

Cooper Snapp

Geminesse Dorsey

Michael Fowler

Ben Greene

William Schneider

Carl Scott

External Tank Thermal Protection System

Myron Pessin

Jim Butler

J. Scott Sparks

Solid Rocket Motor Joint—An Innovative Solution

Paul Bauer

Bruce Steinetz

Ice Detection Prevents Catastrophic Problems

Charles Stevenson

Aerogel-based Insulation System

Charles Stevenson

The Space Shuttle design presented many thermal insulation challenges. The system not only had to perform well, it had to integrate with other subsystems. The Orbiter's surfaces were exposed to exceedingly high temperatures and needed reusable, lightweight, low-cost thermal protection. The vehicle also required low vulnerability to orbital debris and minimal thermal conductivity. NASA decided to bond the Orbiter's thermal protection directly to its aluminum skin, which presented an additional challenge.

The External Tank required insulation to maintain the cryogenic fuels, liquid hydrogen, and liquid oxygen as well as to provide additional structural integrity through launch and after release from the Orbiter.

The challenge and solutions that NASA discovered through tests and flight experience represent innovations that will carry into the next generation of space programs.

Orbiter Thermal Protection System

Throughout the design and development of the Space Shuttle Orbiter Thermal Protection System, NASA overcame many technical challenges to attain a reusable system that could withstand the high-temperature environments of re-entry into Earth's atmosphere.

Theodore von Karman, the dean of American aerodynamicists, wrote in 1956, "Re-entry is perhaps one of the

most difficult problems one can imagine. It is certainly a problem that constitutes a challenge to the best brains working in these domains of modern aerophysics." He was referring to protecting the intercontinental ballistic missile nose cones. Fifteen years later, the shuttle offered considerably greater difficulties. It was vastly larger. Its thermal protection had to be reusable, and this thermal shield demanded both light weight and low cost. The requirement for a fully reusable system meant that new thermal protection

materials would have to be developed, as the technology from the previous Mercury, Gemini, and Apollo flights were only single-mission capable.

Engineers embraced this challenge by developing rigid silica/alumina fibrous materials that could meet the majority of heating environments on windward surfaces of the Orbiter. On the nose cap and wing leading edge, however, the heating was even more extreme. In response, a coated carbon-carbon composite material was developed to


Thermal Protection System Could Take the Heat

Orbiter remained protected during catalytic heating.

While the re-entry surface heating of the Orbiter was predominantly convective, sufficient energy in the shock layer dissociated air molecules and provided the potential for additional heating. As the air molecules broke apart and collided with the surface of the vehicle, they recombined in

an exothermic reaction. Since the surface acted as a catalyst, it was important that the interfacing material/coating have a low propensity to augment the reaction. Atomic recombination influenced NASA's selection of glass-type materials, which have low catalycity and allowed the surface of the

Orbiter to reject a majority of the chemical energy. Engineers performed precise arc jet measurements to quantify this effect over a range of surface temperatures for both oxygen and nitrogen recombination. This resulted in improved confidence in the Thermal Protection System.



form the contours of these structural components. NASA made an exhaustive effort to ensure these materials would operate over a large spectrum of environments during launch, ascent, on-orbit operations, re-entry, and landing.

Environments

During re-entry, the Orbiter's external surface reached extreme temperatures—up to 1,648°C (3,000°F). The Thermal Protection System was designed to provide a smooth, aerodynamic surface while protecting the underlying metal structure from excessive temperature. The loads endured by the system included launch acoustics, aerodynamic loading and associated structural deflections, and on-orbit temperature variations as well as natural environments such as salt fog, wind, and rain. In addition, the Thermal Protection System had to resist pyrotechnic shock loads as the Orbiter separated from the External Tank (ET).

The Thermal Protection System consisted of various materials applied

externally to the outer structural skin of the Orbiter to passively maintain the skin within acceptable temperatures, primarily during the re-entry phase of the mission. During this phase, the Thermal Protection System materials protected the Orbiter's outer skin from exceeding temperatures of 176°C (350°F). In addition, they were reusable for 100 missions with refurbishment and maintenance. These materials performed in temperatures that ranged from -156°C (-250°F) in the cold soak of space to re-entry temperatures that reached nearly 1,648°C (3,000°F). The Thermal Protection System also withstood the forces induced by deflections of the Orbiter airframe as it responded to various external environments.

At the vehicle surface, a boundary layer developed and was designed to be laminar—smooth, nonturbulent fluid flow. However, small gaps and discontinuities on the vehicle surface could cause the flow to transition from laminar to turbulent, thus increasing the overall heating. Therefore, tight fabrication and assembly tolerances were required of the Thermal Protection

System to prevent a transition to turbulent flow early in the flight when heating was at its highest.


Requirements for the Thermal Protection System extended beyond the nominal trajectories. For abort scenarios, the systems had to continue to perform in drastically different environments. These scenarios included: Return-to-Launch Site; Abort Once Around; Transatlantic Abort Landing; and others. Many of these abort scenarios increased heat load to the vehicle and pushed the capabilities of the materials to their limits.

Thermal Protection System Materials


Several types of Thermal Protection System materials were used on the Orbiter. These materials included tiles, advanced flexible reusable surface insulation, reinforced carbon-carbon, and flexible reusable surface insulation. All of these materials used high-emissivity coatings to ensure the maximum rejection of incoming convective heat through radiative heat

Orbiter Tile Placement System Configuration

- Reinforced Carbon-Carbon Coating
- High-temperature Reusable Surface Insulation Tile
- Low-temperature Reusable Surface Insulation Tile
- Advanced Flexible Reusable Surface Insulation Blanket
- Flexible Reusable Surface Insulation Blanket



Orbiter Tile Attachment System High-temperature Reusable Surface Insulation



transfer. Selection was based on the temperature on the vehicle. In areas in which temperatures fell below approximately 1,260°C (2,300°F), NASA used rigid silica tiles or fibrous insulation. At temperatures above that point, the agency used reinforced carbon-carbon.

Tiles

The background to the shuttle's tiles lay in work dating to the early 1960s at Lockheed Missiles & Space Company. A Lockheed patent disclosure provided the first description of a reusable insulation made of ceramic fibers for use as a re-entry vehicle heat shield. In other phased shuttle Thermal Protection System development efforts, ablatives and hot structures were the early competitors. However, tight cost constraints and a strong desire to build the Orbiter with an aluminum airframe pointed toward the innovative, lightweight, and reusable insulation material that could be bonded directly to the airframe skin.

NASA used two categories of Thermal Protection System tiles on the Orbiter—low- and high-temperature

reusable surface insulation. Surface coating constituted the primary difference between these two categories. High-temperature reusable surface insulation tiles used a black borosilicate glass coating that had an emittance value greater than 0.8 and covered areas of the vehicle in which temperatures reached up to 1,260°C (2,300°F). Low-temperature reusable surface insulation tiles contained a white coating with the proper optical properties needed to maintain the appropriate on-orbit temperatures for vehicle thermal control purposes. The low-temperature reusable surface insulation tiles covered areas of the vehicle in which temperatures reached up to 649°C (1,200°F).

The Orbiter used several different types of tiles, depending on thermal requirements. Over the years of the program, the tile composition changed with NASA's improved understanding of thermal conditions. The majority of these tiles, manufactured by Lockheed Missiles & Space Company, were LI-900 (bulk density of 144 kg/m³ [9 pounds/ft³]) and LI-2200 (bulk density of 352 kg/m³ [22 pounds/ft³]).

Fibrous Refractory Composite Insulation tiles helped reduce the overall weight and later replaced the LI-2200 tiles used around door penetrations. Alumina Enhanced Thermal Barrier was used in areas in which small particles would damage fragile tiles. As part of the post-Columbia Return to Flight effort, engineers developed Boeing Rigidized Insulation. Overall, the major improvements included reduced weight, decreased vulnerability to orbital debris, and minimal thermal conductivity.

Orbiter tiles were bonded using strain isolation pads and room-temperature vulcanizing silicone adhesives. The inner mold line of the tile was densified prior to the strain isolation pad bond, which aided in the uniform distribution of the stress concentration loads at the tile-to-strain isolation pad interface. The structure beneath the tile-to-tile gaps was protected by filler bar that prevented gas flow from penetrating into the tile bond line. NASA used gap fillers (prevented hot air intrusion and tile-to-tile contact) in areas of high differential pressures, extreme

aero-acoustic excitations and to passivate over-tolerance step and gap conditions. The structure used for the bonding surface was, for the most part, aluminum; however, several other substrates used included graphite epoxy, beryllium, and titanium.

Design Challenges

Determining the strength properties of the tile-to-strain isolation pad interface was no small feat. The

allowable strength for the interface was approximately 50% less than the LI-900 tile material used on the Orbiter. This reduction was caused by stress concentrations in the reusable surface insulation because of the formation of “stiff spots” in the strain isolation pad by the needling felting process. Accommodating these stiff spots for the more highly loaded tiles was met by locally densifying the underside of the tile. NASA applied

a solution of colloidal silica particles to the non-coated tile underside and baked in an oven at 1,926°C (3,500°F) for 3 hours. The densified layer produced measured about 0.3 cm (0.1 in.) in thickness and increased the weight of a typical 15-by-15-cm (6-by-6-in.) tile by only 27 grams (0.06 pounds). For load distribution, the densified layer served as a structural plate that distributed the concentrated strain isolation pad loads evenly into the weaker, unmodified reusable surface insulation tiles.

NASA faced a greater structural design challenge in the creation of numerous unique tiles. It was necessary to design thousands of these tiles that had compound curves, interfaced with thermal barriers and hatches, and had penetrations for instrumentation and structural access. The overriding challenge was to ensure the strength integrity of the tiles had a probability of tile failure of no greater than $1/10^8$. To accomplish this magnitude of system reliability and still minimize the weight, it was necessary to define the detailed loads and environments on each tile. To verify the integrity of the Thermal Protection System tile design, each tile experienced stresses induced by the following combined sources:

- Substrate or structure out-of-plane displacement
- Aerodynamic loads on the tile
- Tile accelerations due to vibration and acoustics
- Mismatch between tile and structure at installation
- Thermal gradients in the tile
- Residual stress due to tile manufacture
- Substrate in-plane displacement

Other Thermal Protection System Materials? NASA had it Covered.

Flexible Reusable Surface Insulation

White blankets made of coated Nomex® Felt Reusable Surface Insulation protected areas where surface temperatures fell below 371°C (700°F). The blankets were used on the upper payload bay doors, portions of the mid-fuselage, and on the aft fuselage sides.

Advanced Flexible Reusable Surface Insulation

After initial delivery of Columbia to the assembly facility, NASA developed an advanced flexible reusable surface insulation consisting of composite quilted fabric insulation batting sewn between two layers of white fabric. The insulation blankets provided improved producibility and durability, reduced fabrication and installation time and costs, and reduced weight. This insulation replaced the majority of low-temperature reusable surface insulation tiles on two of the shuttles: Discovery and Atlantis.

Following Columbia's seventh flight, the shuttle was modified to replace most of the low-temperature reusable surface insulation tiles on portions of the upper wing. For Endeavour, the advanced flexible reusable surface insulation was directly built into the shuttle.

Additional Materials

NASA used additional materials in other areas of the Orbiter, such as in thermal glass for the windows, Inconel® for the forward Reaction Control System fairings, and elevon seal panels on the upper wing. Engineers employed a combination of white and black pigmented silica cloth for thermal barriers and gap fillers around operable penetrations such as main and nose landing gear doors, egress and ingress flight crew side hatch, umbilical doors, elevon cove, forward Reaction Control System, Reaction Control System thrusters, mid-fuselage vent doors, payload bay doors, rudder/speed brake, and gaps between Thermal Protection System tiles in high differential pressure areas.

Reinforced Carbon-Carbon

The temperature extremes on the nose cap and wing leading edge of the Orbiter required a more sophisticated material that would operate over a large spectrum of environments during launch, ascent, on-orbit operations, re-entry, and landing. Developed by the Vought Corporation, Dallas, Texas, in collaboration with NASA, reinforced carbon-carbon formed the contours of the nose cap and wing leading edge structural components.

Reinforced carbon-carbon is a composite made by curing graphite fabric that has been pre-impregnated with phenolic resin laid up in complex shaped molds. After the parts are rough trimmed, the resin polymer is converted to carbon by pyrolysis—a chemical change brought about by the action of heat. The part is then impregnated with furfuryl alcohol and pyrolyzed multiple times to increase its density with a resultant improvement in its mechanical properties.

Since carbon oxidizes at elevated temperatures, a silicon carbide coating is used to protect the carbon substrate. Any oxidation of the substrate directly affects the strength of the material and, therefore—in the case of the Orbiter—had to be limited as much as possible to ensure high performance over multiple missions. Silicon carbide is formed by converting the outer two plies of the carbon-carbon material through a diffusion coating process, resulting in a stronger coating-to-substrate interlaminar strength.

As a result of the silicon carbide formation, which occurs at temperatures of 1,648°C (3,000°F), craze cracks develop in the coating on cool-down as the carbon substrate


and coating have a different coefficient of thermal expansion. Impregnating the carbon part with tetraethyl orthosilicate and applying a brush-on sealant provides additional protection against oxygen paths to the carbon from the craze cracks.


The tetraethyl orthosilicate is applied via a vacuum impregnation with the intent of filling any remaining porosity within the part. Once the tetraethyl orthosilicate has cured, a silicon dioxide residue coats the pore walls throughout the part, thus inhibiting oxidation. After the tetraethyl orthosilicate process is complete, a sodium silicate sealant is brushed onto the surface of the reinforced carbon-carbon. The sealant fills in the craze cracks and, once cured, forms a glass. The craze cracks close at high temperatures and the sealant will flow

onto the surface; however, since there is sufficient viscosity, the sealant remains on the part. When the reinforced carbon-carbon cools down, the glass fills back into the craze crack.

Why Reinforced Carbon-Carbon?

The functionality of the reinforced carbon-carbon is largely due to its ability to reject heat by external radiation (i.e., giving off heat from surface to the surroundings) and cross-radiation, which is the internal reinforced carbon-carbon heat transfer between the lower and upper structures. Reinforced carbon-carbon has an excellent surface emissivity and can reject heat by radiating to space similar to the other Thermal Protection Systems. It is designed as a shell section with an open interior cavity that promotes cross-radiation.





Since the highest heating is biased toward the lower surface, heat can be cross-radiated to the cooler upper surfaces, thus reducing temperatures of the lower windward surface. Another benefit is that the thermal gradients across the part are minimized.

While reinforced carbon-carbon is designed to withstand high temperatures and maintain its structural shape, the material has a relatively high thermal conductivity so it did not significantly inhibit the heat flow to reach the internal Orbiter wing structure. The metallic attachments that mated the reinforced carbon-carbon to the wing structure were crucial for accommodating the thermal expansion of reinforced carbon-carbon and maintaining a smooth outer mold line of the vehicle. Protecting these attachments and the spar structure itself required internal insulation. Incoflex[®], an insulative batting encased by a thin Inconel[®] foil, protected the metal structural components from the internal cavity radiation environment.

Certification

Prior to the Orbiter's first flight, NASA performed extensive test and analysis to satisfy all requirements related to the natural and induced environments. The space agency accomplished certification of the wing leading edge subsystem for flight by analyses verified with development and qualification tests conducted on full-scale hardware. Engineers performed subscale testing to establish thermal and mechanical properties, while full-scale testing ensured the system performance and provided the necessary data to correlate analytical models. This included a full-scale nose cap test article and twin wing leading edge panel configuration tested through multiple environments (i.e., acoustic/vibration, static loads, and radiant testing). Full-scale testing

ensured that the metallic mechanisms worked in concert with the hot structure as a complete system in addition to meeting the multi-mission requirements.

Reinforced Carbon-Carbon Flight Experience Lessons Learned

While NASA confirmed the fundamental concepts and design sufficiency through the wing leading edge subsystem certification work and early flight test phase of the Space Shuttle Program, the agency also identified design deficiencies. In most cases, modifications rectified those deficiencies. These modifications included addressing the gap heating between the reinforced carbon-carbon and reusable surface insulation to inhibit hot gas flow-through and retrofitting hardware to the wing leading edge subsystem design to account for a substantial increase in the predicted airloads. With increasing design environment maturity, temperature predictions on the attach fittings were significantly lowered, which allowed a design change from steel to titanium and a weight reduction of 136 kg (300 pounds).

Over the 30 years of flight, the shuttle encountered many anomalies that required investigative testing and analysis. Inspections revealed several cracks in the T-seals—i.e., components made of reinforced carbon-carbon that fit between reinforced carbon-carbon panels that allowed for thermal expansion of those components while keeping a smooth outer mold line. The cracks were later found to be caused by convoluted plies from the original layout of the T-seals. NASA corrected the cracking by modifying the manufacturing techniques and implementing additional inspections. In 1993, the agency identified small pinholes that went down to the carbon substrate and were subsequently

traced to a change in maintenance of the launch pad structure. Engineers altered the silica/cement topcoat over the zinc primer such that zinc particles were able to come into contact with the wing leading edge and react with the silicon carbide coating during re-entry, thereby forming pinholes. NASA developed criteria for the pinholes as well as vacuum heat clean and repair methods.

Improved Damage Assessment and Repair With Return to Flight After Columbia Accident

NASA performed rigorous testing and analysis on the Thermal Protection System materials to adequately identify risks and to mitigate failure as much as practical. Engineers developed impact testing, damage-tolerance assessments, and inspection and repair capabilities as part of the Return to Flight effort.

Impact Testing

The greatest lesson learned was that failure of the reinforced carbon-carbon and the catastrophic loss of the vehicle was caused by a large piece of foam debris that was liberated from the ET.

While modifications to the thermal protection foam on the tank reduced the risk of shedding large debris during launch, NASA still expected smaller-sized debris shedding. It was critical that engineers understand the impact of foam shedding on the Orbiter's wing leading edge and tiles. The Southwest Research Institute, San Antonio, Texas, conducted many of these impact tests to understand the important parameters that governed structural failure of reinforced carbon-carbon and tile materials. Additionally, NASA developed finite element modeling capabilities to derive critical-damage thresholds.

Tile Repair—A Critical Capability Was Developed

Prior to the first shuttle launch, NASA recognized the need for a capability to repair tiles on orbit. The loss of a tile during launch due to an improper bond posed the greatest threat. In response, NASA prioritized the development of an ablative material, MA-25S, for repairs of missing or damaged tiles. The biggest obstacle, however, was finding a stable work platform. Thus, NASA cancelled the early repair effort in 1979.

After the Columbia accident in 2003, NASA prioritized tile repair capability. Prior to the Columbia accident, the inspections after every flight revealed damage greater than 2.5 cm (1 in.) in approximately 50 to 100 locations. The original ablative material formed the basis for the repair material developed in the Return to Flight effort.

Some reformulation of MA-25S began in 2003. At that time, NASA changed the

name of the material to Shuttle Tile Ablator, 865 kg/m³ (54 pounds/ft³) (STA-54).

This material decreased the amount of swell during re-entry while maintaining a low enough viscosity to dispense with the extravehicular activity hardware. The material did not harden and would remain workable for approximately 1 hour but still cured within 24 hours in the on-orbit environments.

Simulating a damaged shuttle tile created dust that prevented the STA-54 from penetrating the surface of the tiles. This led to the development of additional materials: a gel cleaning brush that was coated with a sticky silicone substance used to clean tile dust from the repair cavity prior to filling; and primer material that provided a contact surface to which the STA-54 could adhere. Once the primer was cured, the bond strength was stronger than the shuttle tile.



Ground test of Orbiter tile repair.

Finally, NASA performed an on-orbit experiment during STS-123 (2008). Crew member Michael Foreman dispensed STA-54 into several damaged tile specimens. The on-orbit experiment was a success, showing that the material behaved exactly as it had during vacuum dispenses on the ground.

Damage Tolerance Criteria

To make use of the inspection data, NASA developed criteria for critical damage. Damage on reinforced carbon-carbon ranged from spallation (i.e., breaking up or reducing) of the silicon carbide coating to complete penetration of the substrate. Tiles could be gouged by ascent debris to varying depths with a wide variety of cavity shapes. The seriousness of any given damage was highly dependent on local temperature and pressure environments. NASA initiated an extensive Arc Jet test program during Return to Flight activities to characterize the survivability of multiple damage configurations in

different environments. Testing in an Arc Jet facility provided the closest ground simulation for the temperature and chemical constituents of re-entry. Engineers performed numerous tests for both reinforced carbon-carbon and tile to establish damage criteria and verify newly developed thermal math models used for real-time mission support.

Inspection Capability

NASA developed an inspection capability to survey the reinforced carbon-carbon and tile surfaces. This capability provided images to assess any potential impact damages from ascent and orbital debris. A boom with

an imagery sensor package attached to the Shuttle Robotic Arm was used to perform the inspection. The sensor package contained two laser imaging systems and a high-resolution digital camera. Additionally, astronauts residing on the International Space Station (ISS) photographed the entire Orbiter as it executed an aerial maneuver, similar to a backflip, 182 m (600 ft) from the ISS. The crew transmitted photographs to Houston, Texas, where engineers on the ground evaluated the images for any potential damage.

NASA employed an additional detection system to gauge threats from ascent and on-orbit impacts to the wing leading edge. As part of preparing the

Reinforced Carbon-Carbon Repair— Damage Control in the Vacuum of Space

Following the Space Shuttle Columbia accident in 2003, a group of engineers and scientists gathered at Johnson Space Center to discuss concepts for the repair of damaged reinforced carbon-carbon in the weightless vacuum environment of space. Few potential repair materials could withstand the temperatures and pressures on the surface. Of those materials, few were compatible with the space environment and none had been tested in this type of application. Thus, the team developed two repair systems that were made available for contingency use on the next flight.

The first system—Non-Oxide Adhesive Experimental—was designed to repair coating damage or small cracks in reinforced carbon-carbon panels. This pre-ceramic polymer had the consistency of a thick paste. COI Ceramics, Inc., headquartered in San Diego, California,

Orbiter for launch, technicians placed accelerometers on the spar aluminum structure behind the reinforced carbon-carbon panels at the attachment locations. Forty-four sensors across both wings detected accelerations from potential impacts and relayed the data to on-board laptops, which could be transmitted to ground engineers. Using test-correlated dynamic models, engineers assessed suspected impacts for their level of risk based on accelerometer output.



Astronaut Andrew Thomas (left) watches as Charles Camarda tests the reinforced carbon-carbon plug repair (STS-114 [2005]).

developed this system and the NASA repair team slightly modified it to optimize its material properties for use in space. Technicians used a modified commercial caulk gun to apply the material to the damaged wing. The material was spread out over the damage using spatulas similar to commercial trowels. Once dried and cured by the sun, Non-Oxide Adhesive Experimental used the heat of re-entry to convert the material into a ceramic, which protected exposed damage from extreme temperatures and pressures.

For larger damages, a plug repair system protected the reinforced carbon-carbon using a series of thin, flexible composite discs designed to fit securely against the curvature of the surface. Engineers developed 19 geometric shapes, which were flown to provide contingency repair capability. An attach mechanism held the plugs in place. The anchor was made up of a refractory alloy called titanium zirconium molybdenum that was capable of withstanding the 1,648°C (3,000°F) re-entry temperature.

Conclusion

The Orbiter Thermal Protection Systems on the shuttle proved to be effective, with the exception of STS-107 (2003). On that flight, the catastrophic loss was caused by a large piece of foam debris that was liberated from the ET. Advanced materials and coatings were key in enabling the success of the shuttle in high-temperature environments. Experience gathered over many shuttle missions led the Thermal Protection Systems team to

modify and upgrade both design and materials, thus increasing the robustness and safety of these critical systems during the life of the program. Through the tragedy of the Columbia accident, NASA developed new inspection and repair techniques as protective measures to ensure the success and safety of subsequent shuttle missions.



External Tank Thermal Protection System

The amount of Thermal Protection System material on the shuttle's External Tank (ET) could cover an acre. NASA faced major challenges in developing and improving tank-insulating materials and processes for this critical feature. Yet, the space agency's solutions were varied and innovative. These solutions represented a significant advance in understanding the use of Thermal Protection System materials as well as the structures, aerodynamics, and manufacturing processes involved.

The tanks played two major roles during launch: containing and delivering cryogenic propellants to the Space Shuttle Main Engines, and serving as the structural backbone for the attachment of the Orbiter and Solid Rocket Boosters. The Thermal Protection System, composed of spray-on foam and hand-applied insulation and ablator, was applied primarily to the outer surfaces of the tank. It was designed to maintain the quality of the cryogenic propellants, protect the tank structure from ascent heating, prevent the formation of ice (a potential impact debris source), and stabilize tank internal temperature during re-entry into Earth's atmosphere, thus helping to maintain tank structural integrity prior to its breakup within a predicted landing zone.

Basic Configuration

NASA applied two basic types of Thermal Protection System materials to the ET. One type was a low-density, rigid, closed-cell foam. This foam was sprayed on the majority of the tank's "acreage"—larger areas such as the liquid hydrogen and liquid oxygen

tanks as well as the intertank—also referred to as the tank "sidewalls." The other major component was a composite ablator material (a heat shield material designed to burn away) made of silicone resins and cork.

NASA oversaw the development of the closed-cell foam to keep propellants at optimum temperature—liquid hydrogen fuel at -253°C (-423°F) and liquid oxygen oxidizer at -182°C (-296°F)—while preventing a buildup of ice on the outside of the tank, even as the tank remained on the launch pad under the hot Florida sun.

The foam insulation had to be durable enough to endure a 180-day stay at the launch pad, withstand temperatures up to 46°C (115°F) and humidity as high as 100%, and resist sand, salt fog, rain, solar radiation, and even fungus. During launch, the foam had to tolerate temperatures as high as 649°C ($1,200^{\circ}\text{F}$) generated by aerodynamic friction and rocket exhaust. As the tank reentered the atmosphere approximately 30 minutes after launch, the foam helped hold the tank together as temperatures and internal pressurization worked to break it up, allowing the tank to disintegrate safely over a remote ocean location.

Though the foam insulation on the majority of the tank was only about 2.5 cm (1 in.) thick, it added approximately 1,700 kg (3,800 pounds) to the tank's weight. Insulation on the liquid hydrogen tank was somewhat thicker—between 3.8 and 5 cm (1.5 to 2 in.). The foam's density varied with the type, but an average density was 38.4 kg/m^3 (2.4 pounds/ ft^3).

The tank's spray-on foam was a polyurethane material composed of five primary ingredients: an isocyanate and a polyol (both components of the polymeric backbone); a flame


retardant; a surfactant (which controls surface tension and bubble or cell formation); and a catalyst (to enhance the efficiency and speed of the polymeric reaction). The blowing agent—originally chlorofluorocarbon (CFC)-11, then hydrochlorofluorocarbon (HCFC)-141b—created the foam's cellular structure, making millions of tiny bubble-like foam cells.

NASA altered the Thermal Protection System configuration over the course of the Space Shuttle Program; however, by 1995, ET performance requirements led the program to baseline four specially engineered closed-cell foams. The larger sections were covered in polyisocyanurate (an improved version of polyurethane) foam (NCFI 24-124) provided by North Carolina Foam Industries. NCFI 24-124 accounted for 77% of the total foam used on the tank and was sprayed robotically. A similar foam, NCFI 24-57, was sprayed robotically on the aft dome of the liquid hydrogen tank. Stepanfoam® BX-265 was sprayed manually on closeout areas, exterior tank feedlines, and internal tank domes. The tank's ablator, Super-Lightweight Ablator (SLA)-561, was sprayed onto areas subjected to extreme heat, such as brackets and other protuberances, and the exposed, exterior lines that fed the liquid oxygen and liquid hydrogen to the shuttle's main engines. NASA used Product Development Laboratory-1034, a hand-poured foam, for filling odd-shaped cavities.

Application Requirements

Application of the foam, whether automated or hand-sprayed, was designed to meet NASA's requirements for finish, thickness, roughness, density, strength, adhesion, and size and frequency of voids within the foam. The foam was applied in

External Tank Thermal Protection Systems Materials



The External Tank's Thermal Protection System consisted of a number of different foam formulations displayed here. NASA selected materials for their insulating properties, and for their ability to withstand ascent aerodynamic forces.

specially designed, environmentally controlled spray cells and sprayed in several phases, often over a period of several weeks. Prior to spraying, engineers tested the foam's raw material and mechanical properties to ensure the materials met NASA specifications. After the spraying was complete, NASA performed multiple visual inspections of all foam surfaces as well as tests of "witness" specimens in some cases.

More than 90% of the foam was sprayed onto the tank robotically, leaving 10% to be applied by manual spraying or by hand. Most foam was

applied at Lockheed Martin's Michoud Assembly Facility in New Orleans, Louisiana, where the tank was manufactured. Some closeout Thermal Protection System was applied either by hand or manual spraying at the Kennedy Space Center (KSC) in Florida.

Design and Testing

In the early 1970s, NASA developed a spiral "barber pole" Thermal Protection System application technique that was used through the end of the program. This was an early success for the ET Program, but many challenges soon followed.

As the ET was the only expendable part of the shuttle, NASA placed particular emphasis on keeping tank manufacturing costs at a minimum. To achieve this objective, the agency based its original design and manufacturing plans on the use of existing, well-proven materials and processes with a planned evolution to newer products as they became available.

The original baseline Thermal Protection System configuration called for the sprayable Stepanfoam® BX-250 foam (used on the Saturn S-II stage) on the liquid hydrogen sidewalls (acreage)

Solid Rocket Motor Joint—An Innovative Solution


Alliant Techsystems (ATK) Aerospace Systems, in partnership with NASA Glenn Research Center, developed a solution for protecting the temperature-sensitive O-rings used to seal the shuttle reusable solid rocket motor nozzle segments. The use of a carbon fiber material promoted safety and enabled joint assembly in a fraction of the time required by previous processes, with enhanced reproducibility.

The reusable solid rocket motors were fabricated in segments and pinned together incorporating O-ring seals. Similarly, nozzles consisted of multiple components joined and sealed at six joint locations using O-rings. A layer of rubber insulation, referred to as “joint fill” compound, kept the 3,038°C (5,500°F) combustion gases a safe distance away from these seals. In a few instances, however, hot gases breached the compound, leaving soot within the joint. NASA modified the compound installation process and instituted reviews of postflight conditions. Although the modifications proved effective, damage was still possible in the unlikely event that gases breached the compound.

ATK chose an innovative approach through emerging technologies. Rather than attempt to prevent gas intrusion with manually applied rubber fill compound, the heat energy from internal gases would be extracted with a special joint filler and the O-ring seals would be pressurized with the cooled gas.

ATK's solution was based on a pliable, braided form of high-performance carbon material able to withstand harsh temperature environments. The braided design removed most of the thermal energy from the gas and inhibited flow induced by pressure fluctuations. The carbon fiber thermal barrier was easier to install and significantly reduced motor assembly time.

In a rocket environment, carbon fibers withstood temperatures up to 3,816°C (6,900°F). The braided structure and high surface area-to-mass ratio made the barrier an excellent heat exchanger while allowing a restricted yet uniform gas flow. The weave



Using carbon fiber rope instead of rubber insulation in solid rocket motor nozzle joints simplified the joint assembly process and improved shuttle safety margins.

structure allowed it to conform to tolerance assembly conditions. The thermal barrier provided flexibility and resiliency to accommodate joint opening or closing during operation. Upon pressurization, the thermal barrier seated itself in the groove to obstruct hot gas flow from bypassing the barrier.

The carbon fiber solution increased Space Shuttle safety margins. Carbon fibers are suited to a nonoxidizing environment, withstanding high temperatures without experiencing degradation. The barrier provided a temperature drop across a single diameter, reducing gas temperature to O-rings well below acceptable levels. The thermal barrier also kept molten alumina slag—generated during solid fuel burn—from contacting and affecting O-rings.

and forward dome, and SLA-561 (used on the Viking Mars Lander) on the aft dome, intertank, and liquid oxygen tank in the areas of high heating.

In the late 1970s, however, design of the Orbiter tiles advanced to the point where it became apparent that they were susceptible to damage from ice

detaching from the ET. This caused a reassessment of the Thermal Protection System design to prevent the formation of ice anywhere on the tank forward

of the liquid hydrogen tank aft-end structural ring frame. The Orbiter/ice issue drove the requirement to cover the entire tank with Stepanfoam® BX-250, except for the high-heating aft dome, which remained SLA-561. Ice was to be prevented on tank pressurization lines through the use of a heated purge. Certain liquid oxygen feedline brackets, subject to extensive thermal contraction, could not be fully insulated without motion breaking the insulation. Therefore,

NASA accepted ice formation on these brackets as unavoidable.

While attempting to prevent ice buildup on the tank, NASA also worked to characterize both the ablator material and the foams for expected heating rates. NASA worked with Arnold Engineering Development Center in Tennessee to modify its wind tunnel to provide the capability to test foam materials under realistic flight conditions. SLA-561 was tested in the

plasma arc facility at NASA's Ames Research Center in California, which could deliver the required high heating rates. Better understanding of ablation rates and the flow fields around ET protuberances permitted refinement of the Thermal Protection System configuration.

Another unique project was the testing of spray-on foam insulation on a subscale tank, measuring 3 m (10 ft) in diameter, in the environmental hanger at Eglin Air Force Base, Florida. The insulated tank was filled with liquid nitrogen and subjected to various rain, wind, humidity, and temperature conditions to determine the rate of ice growth. These data were then converted to a computer program known as Surfice, which was used at KSC to predict whether unacceptable ice would form prior to launch.

To provide information on application techniques, the agency ran cryogenic flexure tests that verified substrate adhesion and strength as well as crush tests on the Thermal Protection System materials.

In a continuous search for optimum Thermal Protection System performance, NASA—still in the Thermal Protection System design and testing phase—decided to use Chemical Products Research (CPR)-421, a commercial foam insulation with good high-heating capability. Lockheed Martin developed a sprayable Thermal Protection System to apply to tank sidewalls and aft dome. Application needed a relative humidity of less than 30%, which resulted in the addition of a chemical dryer at Michoud. Also, the tank wall had to be heated to 60°C (140°F). This required passing hot gas through the tank while it was being rotated for the “barber pole” foam application mode.



A secondary function of the Thermal Protection System was to stabilize tank internal temperature during re-entry into Earth's atmosphere, thus helping to maintain tank structural integrity prior to its breakup over a remote ocean location.



The key to the External Tank's foam Thermal Protection System insulating properties was its cellular structure, creating millions of tiny bubble-like foam cells. The sprayed foam (NCFI 24-124) can be seen here after application to an area of the tank's aluminum “acreage,” consisting of the liquid oxygen tank, liquid hydrogen tank, and intertank.

Ice Detection Prevents Catastrophic Problems

NASA had a potentially catastrophic problem with ice that formed on the cryogenic-filled Space Shuttle External Tank. Falling ice could have struck and damaged the crew compartment windows, reinforced carbon-carbon panels on the wing leading edge of the Orbiter, or its thermal protection tiles, thus placing the crew and vehicle at risk.

Kennedy Space Center and the US Army Tank Automotive and Armaments Research, Development and Engineering Center confirmed that a proof-of-concept system, tested by MacDonald, Dettwiler and Associates Ltd. of Canada, offered potential to support cryogenic tanking tests and ice debris team inspections on the launch

pad. NASA and its partners initiated a program to develop a system capable of detecting ice on the External Tank spray-on foam insulation surfaces. This system was calibrated for those surfaces and used an infrared strobe, a focal plane sensor array, and a filter wheel to collect successive images over a number of sub-bands. The camera processed the images to determine whether ice was present, and it also computed ice thickness. The system was housed in nitrogen-purged enclosures that were mounted on a two-wheeled portable cart. It was successfully applied to the inspection of the External Tank on STS-116 (2006), where the camera detected thin ice/frost layers on two umbilical connections.



Robert Speece, NASA engineer, is shown operating the ice detection system at the pad, prior to shuttle launch.

The system can be used to detect ice on any surface. It can also be used to detect the presence of water.

First Flight Approaches

As the Space Shuttle Program moved toward the first shuttle flight in 1981, NASA faced another challenge. Approximately 37 m² (400 ft²) of ablator became debonded from the tank's aluminum surface the first time a tank was loaded with liquid hydrogen. While the failure analysis was inconclusive, it appeared that the production team had tried to bond too large an area and did not get the ablator panels under the required vacuum before the adhesive pot life ran out. Technicians at Michoud Assembly Facility reworked the application process for the ET at their facility and the first tank at KSC.

Following the ablator bonding problem, NASA intensified its analysis of the ablator/aluminum bond line. This analysis showed that the higher

coefficient of thermal expansion of the ablator binder, as compared to the aluminum, would cause the ablator to shrink. This would introduce biaxial tension in the ablator and corresponding shear forces at the bond line near any edges, discontinuities, or cracks. Then, when the tank was pressurized, tank expansion from pressure would compound this shear force, possibly causing the bond line to fail. NASA decided to pre-pressurize the liquid hydrogen tank with helium gas prior to filling the tank for launch—and to pressures higher than flight pressures—to stretch the ablator when it was warm and elastic.

Because early test data showed the tank insulation could be adversely affected by ultraviolet light, NASA painted the first several tanks white, using a fire-retardant latex paint. Exposure

testing of foam samples on the roof of the Michoud Assembly Facility, however, showed the damage to be so shallow that it was insignificant. NASA decided not to paint the tanks, resulting in a weight savings of about 260 kg (580 pounds), lowered labor costs, and the introduction of the “orange” tank.

Environmental Challenges


Knowledge of toxic properties and environmental contaminations increased over the 30 years of the Space Shuttle Program. Federal laws reflected these changes. For instance, ozone-depleting substances, including some Freon® compounds, reduced the protecting atmospheric ozone layer. NASA worked with its contractors to reduce both toxicity and environmental consequences for the cooling agents and the foam compounds.

During the 1990s, the University of Utah published data showing that CPR-421 was potentially toxic. Based on this analysis, Chemical Products Research withdrew CPR-421 from the market. NASA's ET office had Chemical Products Research reformulate this foam, with the new product identified as CPR-488.

New challenges arose related to emerging environmental policies that necessitated changes to Thermal Protection System foam formulations. In 1987, the United States adopted the Montreal Protocol on Substances that Deplete the Ozone Layer, which provided for the eventual international elimination of ozone-depleting substances. The United States implemented the protocol by regulations under the Clean Air Act. Ozone-depleting substances, including CFC-11—the Freon® blowing agent used in the production of the Thermal Protection System sprayable foams for the tanks—were scheduled to be phased out of production. After the phaseout, CFC-11 would only be available for such uses through a rigorous exemption process.

To prepare for the upcoming obsolescence of the foam blowing agent, Marshall Space Flight Center (MSFC) along with Lockheed Martin tracked and mitigated the effect of emerging environmental regulations. After extensive research and testing of potential substitutes, NASA proposed that HCFC-141b replace the CFC-11 blowing agent. NASA continued to use stockpiled supplies of CFC-11-blown foam until the HCFC-141b foam was certified for tank use and phased in beginning in 1996.

NASA undertook the development and qualification of a foam to be phased in as a replacement for the tank



The foam's approximately 2.5-cm (1-in.) thickness borders the circumferential flange that joins the intertank with the liquid hydrogen tank. The ribbed area is the intertank, that, like the liquid oxygen tank in the background and the liquid hydrogen tank in the foreground, was robotically sprayed with NCFI 24-124 foam. The flange would later be hand-sprayed with Stepanfoam® BX-265. The liquid oxygen feedline at the top of the tank and a feedline bracket have been hand-sprayed with BX-265 foam.



A technician at NASA's Michoud Assembly Facility sprays the flange that connects the intertank and liquid hydrogen tank. Stepanfoam® BX-265 was sprayed manually on closeout areas, exterior tank feedlines, internal tank domes, closeout areas of mating External Tank subcomponent surfaces, and small subcomponents.

sidewall foam, CPR-488. North Carolina Foam Industries reformulated CPR-488 and developed a new product.

As part of qualifying this new product, Lockheed Martin, Wyle Laboratories,

and MSFC developed an environmental test. This test used a flat aluminum plate machined to match aft dome stress levels. The plate was attached to a cryostat filled with liquid helium and then strained with hydraulic jacks

to the flight biaxial stress levels. Radiant heat lamps were installed to match the radiant heating from the solid rocket motor plumes, and an acoustic horn blasted the test. This simulated the aft dome ascent environment as well as possible. The test results indicated the need to spray ablator on the aft dome. To provide the capability to spray the ablator, personnel at Michoud Assembly Facility built two spray cells, with an additional cell to clean and prime the liquid hydrogen tank before ablator application.

To save the weight of this ablator and its associated cost, NASA had North Carolina Foam Industries develop a foam adequate for the aft dome environment without ablator. The foam was phased in on the aft dome, flying first on Space Transportation System (STS)-79 in 1996. The first usage of the new foam on the tank sidewalls was phased in over three tanks starting with STS-85 in August 1997.

Environmental Protection Agency regulations also required NASA to replace Stepanfoam® BX-250, which


was sprayed manually—with a CFC-11 blowing agent—on the tank's "closeout" areas. During STS-108 (2001), Stepanfoam® BX-265—with HCFC-141b as its blowing agent—first flew as a replacement for BX-250. BX-250 continued to be flown in certain applications as BX-265 was phased into the manufacturing process.

The use of HCFC-141b as a foam blowing agent, however, was also problematic. It was classified as a Class II ozone-depleting substance and was subject to phaseout under the

Aerogel-based Insulation System Precluded Hazardous Ice Formation

During the STS-114 (2005) tanking test, the External Tank Gaseous Hydrogen Vent Arm Umbilical Quick Disconnect formed ice and produced liquid nitrogen/air. The phenomenon was repeated during subsequent testing and launch. For the shuttle, ice presented a debris hazard to the Orbiter Thermal Protection System and was unacceptable at this umbilical location. The production of uncontrolled liquid nitrogen/air presented a hazard to the shuttle, launch pad, and ground support equipment.

NASA incorporated a fix into the existing design to preclude ice formation and the uncontrolled production of liquid nitrogen/air. The resolution was



Testing of gaseous hydrogen vent arm umbilical disconnect equipment at Kennedy Space Center.

accomplished with two changes to the umbilical purge shroud. First, the space agency improved the shroud purge gas flow to obtain the desired purge cavity gas concentrations. Second, technicians wrapped multiple layers of aerogel blanket material directly onto the quick disconnect metal surfaces within the purged shroud cavity.

NASA tested the design modifications at the Kennedy Space Center Cryo Test Lab. Tests showed that the outer surface of the

shroud was maintained above freezing with no ice formation and that no nitrogen penetrated into the shroud purge cavity. NASA used the modified design on STS-121 (2006) and all subsequent flights.

Aerogel insulation is a viable alternative to the current technology for quick disconnect shrouds purged with helium or nitrogen to preclude the formation of ice and liquid nitrogen/air. In most cases, aerogel insulation eliminates the need for active purge systems.

Clean Air Act effective January 2003. NASA was granted exemptions permitting the use of HCFC-141b in foams for specific shuttle applications. These exemptions applied until the end of the program.

Post-Columbia Accident Advances in Thermal Protection

Following the loss of Space Shuttle Columbia in 2003, NASA undertook the redesign of some tank components to reduce the risk of ice and foam debris coming off the tank. These hardware changes drove the need to improve the application of Thermal Protection System foam that served as an integral part of the components' function. The major hardware addressed included the ET/Orbiter attach bipod closeout, protuberance air load ramps, ice frost ramps, and the liquid hydrogen tank-to-intertank flange area.

The ET bipod attached the Orbiter to the tank. The redesign removed the foam ramps that had covered the bipod attach fittings, and which had been designed to prevent the formation of ice when the ET was filled with cold liquid hydrogen and liquid oxygen on the launch pad. This left the majority of each fitting exposed. NASA installed heaters as part of the bipod configuration to prevent ice formation on the exposed fittings.

NASA developed a multistep process to improve the manual bipod Thermal Protection System spray technique. Validation of this process was accomplished on a combination of high-fidelity mock-ups and a full-scale ET test article in a production

environment. Wind tunnel tests demonstrated Thermal Protection System closeout capability to withstand maximum aerodynamic loads without generating debris.


The ET protuberance air load ramps were manually sprayed wedge-shaped layers of insulating foam insulation along the pressurization lines and cable tray on the side of the tank. They were designed as a safety precaution to protect the tank's cable trays and pressurization lines from airflow that could potentially cause instability in these attached components. Foam loss from the ramps during ascent, however, drove NASA to remove them from

the tank. This required extensive engineering. NASA created enhanced structural dynamics math models to better define the characteristics of this area of the tank and performed numerous wind tunnels tests.

The ET fuel tank Main Propulsion System pressurization lines and cable trays were attached along the length of the tank at multiple locations by metal support brackets. These were protected from forming ice and frost during tanking operations by foam protuberances called ice frost ramps. The feedline bracket configuration had the potential for foam and ice debris loss. Redesign changes were

External Components Redesign

Original Configuration



After the Columbia accident, NASA implemented a number of improvements to External Tank components and related Thermal Protection System elements. One such measure was the redesign of the Orbiter/External Tank attach bipod fitting mechanism, which included a meticulous reworking of the attach fitting Thermal Protection System configuration.





In what used to be a one-person operation, a team of technicians at NASA's Michoud Assembly Facility prepares to hand-spray BX-250 foam on the bipod attach fittings. The videographer (standing) records the process for later review and verification. A quality control specialist (left) witnesses the operation, while two spray technicians make preparations.

incorporated into the 17 ice frost ramps on the liquid hydrogen tank to reduce foam loss. BX-265 manual spray foam replaced foam in the ramps' closeout areas to reduce debonding and cracking.

The NASA/Lockheed Martin team also developed an enhanced three-part procedure to improve the Thermal Protection System closeout process on the liquid hydrogen tank-to-intertank flange area.

In all post-Columbia Thermal Protection System enhancement efforts, NASA modified process controls to ensure that defects were more tightly kept within the design envelope. The space agency simplified application techniques and spelled out instructions in more detail, and technicians had the opportunity to practice their application skills on high-fidelity component models. MSFC and Lockheed Martin also developed an electronic database to store information for each spray. New application certification requirements were added. Improvements included the forward bellows heater, the liquid oxygen feedlines, and titanium brackets. Improved imagery analysis and probabilistic risk assessments also allowed NASA to better track and predict foam loss. Thermal protection debris could never be completely eliminated, but NASA had addressed a complex and unprecedented set of problems with determination and innovation.



NASA decided to delete the tank's protuberance air load ramps and implement design changes to the 17 ice frost ramps on the liquid hydrogen tank. Both these measures required adjustments in the components' Thermal Protection System configuration and application processes. Materials and techniques were also altered to improve the Thermal Protection System closeout of the flange joining the liquid hydrogen tank with the intertank.



Materials and Manufacturing

Introduction

Gail Chapline

Nondestructive Testing Innovations

Willard Castner

Patricia Howell

James Walker

Friction Stir Welding Advancements

Robert Ding

Jim Butler

Characterization of Materials in the Hydrogen Environment

Jon Frandsen

Jonathan Burkholder

Gregory Swanson

Space Environment: It's More Than a Vacuum

Lubert Leger

Steven Koontz

Chemical Fingerprinting

Michael Killpack

Environmental Assurance

Anne Meinhold

Unprecedented Accomplishments in the Use of Aluminum-Lithium Alloy

Preston McGill

Jim Butler

Myron Pessin

Orbiter Payload Bay Door

Lubert Leger

Ivan Spiker

To build a spacecraft, we must begin with materials. Sometimes the material choice is the solution. Other times, the design must accommodate the limitations of materials properties. The design of the Space Shuttle systems encountered many material challenges, such as weight savings, reusability, and operating in the space environment. NASA also faced manufacturing challenges, such as evolving federal regulations, the limited production of the systems, and maintaining flight certification. These constraints drove many innovative materials solutions. Innovations such as large composite payload bay doors, nondestructive materials evaluation, the super lightweight tank, and the understanding of hydrogen effects on materials were pathfinders used in today's industry. In addition, there were materials innovations in engineering testing, flight analysis, and manufacturing processes. In many areas, materials innovations overcame launch, landing, and low-Earth orbit operational challenges as well as environmental challenges, both in space and on Earth.



Nondestructive Testing Innovations

Have you ever selected a piece of fruit based on its appearance or squeezed it for that certain feel? Of course you have. We all have. In a sense, you performed a nondestructive test. Actually, we perform nondestructive testing every day. We visually examine or evaluate the things we use and buy to see whether they are suitable for their purpose. In most cases, we give the item just a cursory glance or squeeze; however, in some cases, we give it a conscious and detailed examination. We don't think of these routine examinations as nondestructive tests, but they are, and they give us a sense of what nondestructive testing is about.

Nondestructive testing is defined as the inspection or examination of materials, parts, and structures to determine their integrity and future usefulness without compromising or affecting their usefulness. The most fundamental nondestructive test of all is visual

inspection. In the industrial world, visual examination can be quite formal, with complex visual aids, pass/fail criteria, training requirements, and written procedures.

Nondestructive testing depends on incident or input energy that interacts with the material or part being examined. The incident or input energy can be modified by reflection from interaction within or transmission through the material or part. The process of detection and interpretation of the modified energy is how nondestructive testing provides knowledge about the material or part. Tests range from the simple detection and interpretation of reflected visible light by the human eye (visual examination) to the complex electronic detection and mathematical reconstruction of through-transmitted x-radiation (computerized axial tomography [CAT] scan). From a nondestructive testing perspective, the similarity between the simple visual examination and the complex CAT scan is the input energy (visible light vs. x-rays) and the modified energy

(detected by the human eye vs. an electronic x-ray detector).

Nondestructive testing is a routine part of a spacecraft's life cycle. For the reusable shuttle, nondestructive testing began during the manufacturing and test phases and was applied throughout its service life. NASA performed many such nondestructive tests on the shuttle vehicles and developed most nondestructive testing innovations in response to shuttle problems.

Quantitative Nondestructive Testing of Fatigue Cracks

One of the most significant nondestructive testing innovations was quantifying the flaw sizes that conventional nondestructive testing methods could reliably detect. NASA used artificially induced fatigue cracks to make the determination because such flaws were relatively easy to grow and control, hard to detect, and tended to bound the population of flaws of interest. The need to quantify the reliably detectable crack sizes was



Two examples of the most basic nondestructive testing:

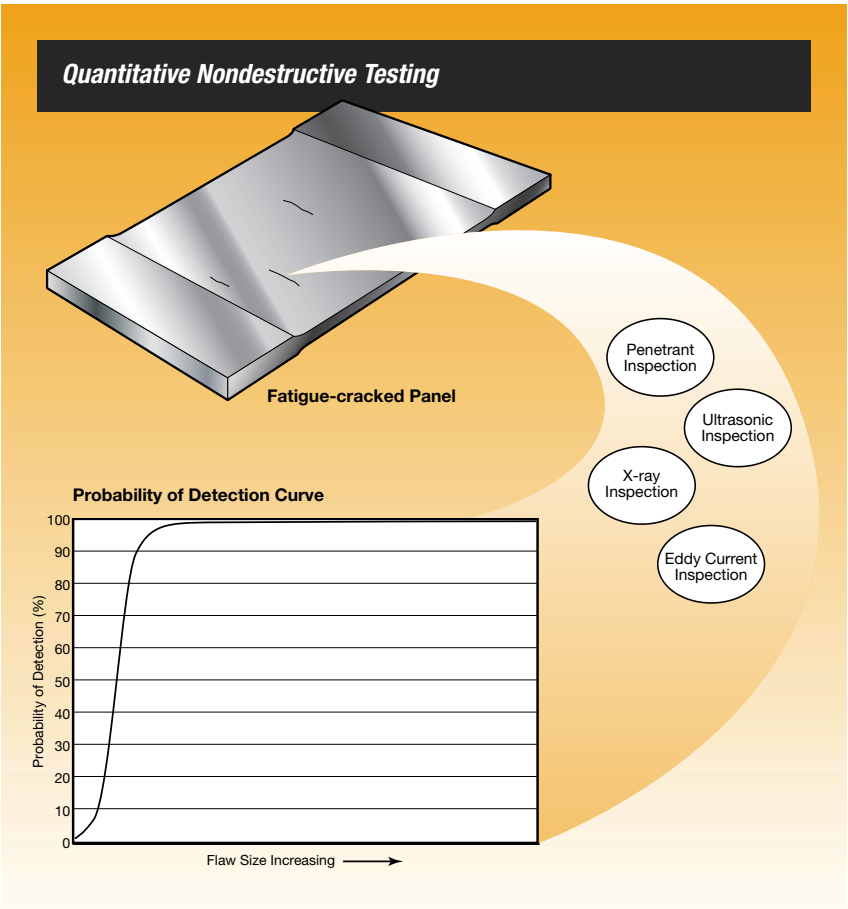
Left, a gardener checks ripening vegetables. Right, Astronaut Eileen Collins, STS-114 (2005) mission commander, looks closely at a reinforced carbon-carbon panel on one of the wings of the Space Shuttle Atlantis in the Orbiter Processing Facility at Kennedy Space Center (KSC). Collins and the other crew members were at KSC to take part in hands-on equipment and Orbiter familiarization.



mandated by a fracture control interest in having confidence in the starting crack size that could be used in fracture and life calculations. Although there was no innovation of any specific nondestructive testing method, quantifying—in a statistical way—the reliably detectable crack sizes associated with the conventional nondestructive evaluation methods was innovative and led the way to the adoption of similar quantitative nondestructive evaluation practices in other industries.

The quantification of nondestructive testing methods is commonly referred to today as probability of detection. The Space Shuttle Program developed some of the earliest data for the penetrant, x-ray, ultrasonic, and eddy current nondestructive testing methods—the principal nondestructive testing methods used to inspect shuttle components during manufacturing. Data showed that inspectors certified to aerospace inspection standards could, on average, perform to a certain probability of detection level defined as standard nondestructive evaluation.

Beyond standard nondestructive evaluation, NASA introduced a special nondestructive evaluation level of probability of detection wherein the detection of cracks smaller than the standard sizes had to be demonstrated by test. Engineers fabricated fatigue-cracked specimens that were used over many years to certify and recertify, by test, the inspectors and their nondestructive evaluation processes to the smaller, special nondestructive evaluation crack size. The size of the fatigue cracks in the specimens was targeted to be a surface-breaking semicircular crack 0.127 cm (0.050 in.) long by 0.063 cm (0.025 in.) deep, a size that was significantly smaller than the standard nondestructive evaluation crack size of 0.381 cm (0.150 in.) long by 0.19 cm (0.075 in.) deep.



The special probability of detection specimen sets typically consisted of 29 randomly distributed cracks of approximately the same size. By detecting all 29 cracks, the inspector and the specific nondestructive evaluation process were considered capable of detecting the crack size to a 90% probability of detection with 95% confidence.

Nondestructive Testing of Thermal Protection System Tiles

The development of Thermal Protection System tiles was one of the most unique and difficult developments of the program. Because of this material's "unknowns," the tile attachment scheme, and their extremely fragile

nature, NASA examined a number of nondestructive testing methods.

Acoustic Emission Monitoring


Late in the development of the shuttle Thermal Protection System and just before the first shuttle launch, NASA encountered a major problem with the attachment of the tiles to the Orbiter's exterior skin. The bond strength of the tile system was lower than the already-low strength of the tile material, and this was not accounted for in the design. The low bond strength was due to stress concentrations at the tile-to-strain isolation pad bond line interface. A Nomex® felt strain isolation pad was bonded between each tile and the Orbiter skin to minimize the



lateral strain input to the tile from the aluminum skin. These stress concentrations led to early and progressive failures of the tile material at the tile-to-strain isolation pad bond line interface when the tile was loaded.

To determine whether low bond strengths existed, engineers resorted to proof testing for each tile. This required thousands of individual tile proof tests prior to first flight. Space Shuttle Columbia (Space Transportation System [STS]-1) was at Kennedy Space Center being readied for first flight when NASA decided that proof testing was necessary. Since proof testing was not necessarily nondestructive and tiles could be damaged by the test, NASA sought a means of monitoring potential damage; acoustic emission nondestructive testing was an obvious choice. The acoustic signatures of a low bond strength tile or a tile damaged during proof test were determined through laboratory proof testing of full-size tile arrays.

To say that the development and implementation of acoustic emission monitoring during tile proof testing was done on a crash basis would be an understatement. The fast pace was dictated by a program that was already behind schedule, and the tile bond strength problem threatened significant additional delay. At the height of the effort, 18 acoustic emission systems with fully trained three-person crews were in operation 24 hours a day, 7 days a week. The effort was the largest single concentration of acoustic emission equipment at a single job site. As often happens with such problems, where one solution can be overtaken and replaced by another, a tile densification design fix for the low-strength bond was found and implemented prior to first flight, thus obviating the need for continued



acoustic emission monitoring. By the time the acoustic emission monitoring was phased out, NASA had performed 20,000 acoustic emission monitored proof tests.

Sonic Velocity Testing


Another early shuttle nondestructive testing innovation was the use of an ultrasonic test technique to ensure that the Thermal Protection System tiles were structurally sound prior to installation. Evaluation of pulse or sonic velocity tests showed a velocity relationship with respect to

both tile density and strength. These measurements could be used as a quality-control tool to screen tiles for low density and low strength and could also determine the orientation of the tile.

The sonic velocity technique input a short-duration mechanical impulse into the tile. A transmitting transducer and a receiving transducer, placed on opposite sides of the tile, measured the pulse's transit time through the tile. For the Lockheed-provided tile material, LI-900 (with bulk density of 144 kg/m³ [9 pounds/ft³]), the average through-the-thickness sonic velocity was on the

Sonic Velocity Testing of Tiles at Kennedy Space Center Thermal Protection System Facility

The speed of sound through the tile is related to density and strength.



order of 640 m/sec (2,100 ft/sec), and the through-the-thickness flat-wise tensile strength was on the order of 1.69 kg/cm² (24 pounds/in²). The LI-900 acceptance criterion for sonic velocity was set at 518 m/sec (1,700 ft/sec), which corresponded to a minimum strength of 0.91 kg/cm² (13 pounds/in²). Sonic velocity testing was phased out in the early 1990s.

Post-Columbia Accident Nondestructive Testing of External Tank

A consequence of the Columbia (STS-107) accident in 2003 was the development of several nondestructive innovations, including terahertz imaging and backscatter radiography of External Tank foam and thermography of the reinforced carbon-carbon—both on orbit and on the ground—during vehicle turnaround. The loss of foam, reinforced carbon-carbon impact damage, and on-orbit inspection of Thermal Protection System damage were all problems that could be mitigated to some extent through the application of nondestructive testing methods.

Nondestructive Testing of External Tank Spray-on Foam Insulation

Prior to the Columbia accident, no nondestructive testing methods were available for External Tank foam inspection, although NASA pursued development efforts from the early 1980s until the early 1990s. The foam was effectively a collection of small air-filled bubbles with thin polyurethane membranes, making the foam a thermal and electrical insulator with very high acoustic attenuation. Due to these properties, it was not feasible to inspect the foam with conventional methods such as eddy current, ultrasonics, or thermography. In addition, since the foam was considered nonstructural, problems of delaminations occurring during foam application and foam popping off (“popcorning”) during ascent were considered manageable through process control.

After the Columbia accident, NASA focused on developing nondestructive testing methods for finding voids and delaminations in the thick, hand-sprayed foam applications around protuberances and closeout

areas. The loss of foam applied to the large areas of the tank was not as much of concern because the automated acreage spray-on process was better controlled, making it more unlikely to come off. In the event it did come off, the pieces would likely be small because acreage foam was relatively thin. NASA’s intense focus resulted in the development and implementation of two methods for foam inspection—terahertz imaging and backscatter radiography—that represented new and unique application of nondestructive inspection methods.


Terahertz Imaging

Terahertz imaging is a method that operates in the terahertz region of the electromagnetic spectrum between microwave frequencies and far-infrared frequencies. Low-density hydrocarbon materials like External Tank foam were relatively transparent to terahertz radiation. Terahertz imaging used a pulser to transmit energy into a structure and a receiver to record the energy reflected off the substrate or internal defects. As the signal traveled through the structure, its basic wave



Terahertz Imaging System

This system uses high-frequency electromagnetic pulses.



properties were altered by the attenuation of the material and any internal defects. An image was made by scanning the pulser/receiver combination over the foam surface and displaying the received signal.

Probability of detection studies of inserted artificial voids showed around 90% detection of the larger voids in simple geometries, but less than 90% detection in the more-complicated geometries of voids around protrusions. Further refinements showed that delaminations were particularly difficult to detect. The detection threshold for a 2.54-cm- (1-in.)-diameter laminar defect was found to be a height of 0.508 cm (0.2 in.), essentially meaning delaminations could not be detected. The terahertz inspection method was used for engineering evaluation, and any defects found were dealt with by an engineering review process.

Backscatter Radiography


Backscatter radiography uses a conventional industrial x-ray tube to generate a collimated beam of x-rays

that is scanned over the test object. The backscattering of x-rays results from the Compton effect—or scattering—in which absorption of the incident or primary x-rays by the atoms of the

Backscatter X-ray Imaging System



Insulating foam covers the External Tank.



An irradiated column of foam that has voids produces less backscattered x-rays than a void-free column of foam.



test material are reradiated at a lower energy as secondary x-rays in all directions. The reradiated or backscattered x-rays were collected in collimated radiation detectors mounted around the x-ray source. Voids or defects in the test material were imaged in backscatter radiography in the same manner as they were in conventional through-transmission radiography. Imaging of voids or defects depended on less absorbing material and less backscattered x-rays from the void.

Since only the backscattered x-rays were collected, the technique was single sided and suited for foam inspection. The foam was well suited for backscatter radiography since Compton scattering is greater from low atomic number materials. The technique was more sensitive to near surface voids but was unable to detect delaminations. Like terahertz imaging, backscatter radiography was used for engineering evaluation, and defects found were dealt with by an engineering review process.

Nondestructive Testing of Reinforced Carbon-Carbon System Components

A recommendation of the Columbia Accident Investigation Board stated: "Develop and implement a comprehensive inspection plan to determine the structural integrity of all Reinforced Carbon-Carbon (RCC) system components. This inspection plan should take advantage of advanced non-destructive inspection technology." To comply with this recommendation, NASA investigated advanced inspection technology for inspection of the reinforced carbon-carbon leading edge panels during ground turnarounds and while on orbit.

Ground Turnaround Thermography

NASA selected infrared flash thermography as the method to determine the structural integrity of the reinforced carbon-carbon components. Thermography was a fast, noncontacting, one-sided application that was easy to implement in the Orbiter's servicing environment.



Infrared thermography inspection of the Orbiter nose captured at the instant of the xenon lamp flash. Kennedy Space Center Orbiter Processing Facility.

The Thermographic Inspection System was an active infrared flash thermography system. Thermographic inspection examined and recorded the surface temperature transients of the test article after application of a short-duration heat pulse. The rate of heat transfer away from the test article surface depended on the thermal diffusivity of the material and the uniformity and integrity of the test material. Defects in the material would retard the heat flow away from the surface, thus producing surface temperature differentials that were reflective of the uniformity of the material and its defect content. A defect-free material would uniformly

transfer heat into the underlying material, and the surface temperature would appear the same over the entire test surface; however, a delamination would prevent or significantly retard heat flow across the gap created by the delamination, resulting in more-local heat retention and higher surface temperature in comparison to the material surrounding the delamination. Temperature differences were detected by the infrared camera, which provided visual images of the defects. Electronic signals were processed and enhanced for easier interpretation. The heat pulse was provided by flashing xenon lamps in a hooded arrangement that excluded ambient light. The infrared camera was transported along a floor-mounted rail system in the Orbiter Processing Facility for the leading edge panel inspections, allowing full and secure access to all of the leading edge surfaces. After the transport cart was positioned, the camera was positioned manually via a grid system that allowed the same areas to be compared from flight to flight.

The thermography system was validated on specimens containing flat bottom holes of different diameters and depths. Validation testing confirmed the ability of the flash thermography system to detect the size holes that needed to be detected.

After the first Return to Flight mission—STS-114 (2005)—the postflight thermography inspection discovered a suspicious indication in the joggle area of a panel. Subsequent investigation showed that the indication was a delamination. This discovery set in motion an intense focus on joggle-area delaminations and their characterization and consequence. Many months of further tests, development, and refinement of the thermography methodology



determined that critical delaminations would be detected and sized by flash thermography and provided the basis for flightworthiness.

On-orbit Thermography


The success of infrared thermography for ground-based turnaround inspection of the wing leading edge panels and the extensive use of thermography during Return to Flight impact testing made it the choice for on-orbit inspection of the leading edge reinforced carbon-carbon material. A thermal gradient through the material must exist to detect subsurface reinforced carbon-carbon damage with infrared thermography. A series of ground tests demonstrated that sunlight or solar heating and shadowing could be used to generate the necessary thermal gradient, which significantly simplified the camera development task.

With the feasibility of on-orbit thermography demonstrated and with the spaceflight limitations on weight and power taken into account, NASA selected a commercial off-the-shelf microbolometer camera for modification and development into a space-qualified infrared camera for inspecting the reinforced carbon-carbon for impact damage while on orbit.


The extravehicular activity infrared camera operated successfully on its three flights. Two reinforced carbon-carbon test panels with simulated damage were flown and inspected on STS-121 (2006). The intentional impact damage in one panel and the flat bottom holes in the other panel were clearly imaged. Engineers also performed a similar on-orbit test on two other intentionally damaged reinforced carbon-carbon test panels during a space station extravehicular activity with the

On-orbit Thermography

Processed infrared images of reinforced carbon-carbon test panels.



Astronaut Thomas Reiter mounting pre-damaged reinforced carbon-carbon test panels on the International Space Station during STS-121 (2006).



Extravehicular activity infrared flight camera.



same result of clearly imaging the damage. The end result of these efforts was a mature nondestructive inspection technique that was transitioned and demonstrated as an on-orbit nondestructive inspection technique.

Additional Nondestructive Testing

Most nondestructive testing innovations resulted from problems that the shuttle encountered over the years, where nondestructive testing provided all or part of the solution. Other solutions worth mentioning include: ultrasonic extensometer measurements of critical shuttle bolt tensioning; terahertz imaging of corrosion under tiles; phased array ultrasonic testing of the External Tank friction stir welds and the shuttle crawler-transporter shoes; thermographic leak detection of the main engine nozzle; digital radiography of Columbia debris; surface replication of flow liner cracks; and the on-board wing leading edge health monitoring impact system.

Friction Stir Welding Advancements

NASA invents welding fixture.



Friction stir welding units, featuring auto-adjustable pin tools, welded External Tank barrel sections at NASA's Michoud Assembly Facility in New Orleans, Louisiana. The units measured 8.4 m (27.5 ft) in diameter and approximately 7.6 m (25 ft) tall to accommodate the largest barrel sections.

In the mid 1990s, NASA pursued the implementation of friction stir welding technology—a process developed by The Welding Institute of Cambridge, England—to improve External Tank welds. This effort led to the invention of an auto-adjustable welding pin tool adopted by the Space Shuttle Program, the Ares Program (NASA-developed heavy launch vehicles), and industry.

Standard fusion-welding techniques rely on torch-generated heat to melt and join the metal. Friction stir welding does not melt the metal. Instead, it uses a rotating pin and “shoulder” to generate friction, stir the metal together, and forge a bond. This process results in welds with mechanical properties superior to fusion welds.

Standard friction stir welding technology has drawbacks, however; namely, a non-adjustable pin tool that leaves a “keyhole” at the end of a circular weld and the inability to automatically adjust the pin length for materials of varying thickness. NASA's implementation of friction stir welding for the External Tank resulted in the invention and patenting of an auto-adjustable pin tool that automatically retracts and extends in and out of the shoulder. This feature provides the capability to make 360-degree welds without leaving a keyhole, and to weld varying thicknesses.

During 2002-2003, NASA and the External Tank prime contractor, Lockheed Martin, implemented auto-adjustable pin tool friction stir welding for liquid hydrogen and liquid oxygen tank longitudinal welds. Since that time, these friction stir welds have been virtually defect-free. NASA's invention was being used to weld Ares upper-stage cryogenic hardware. It has also been adopted by industry and is being used in the manufacturing of aerospace and aircraft frames.



Characterization of Materials in the Hydrogen Environment

From the humid, corrosion-friendly atmosphere of Kennedy Space Center, to the extreme heat of ascent, to the cold vacuum of space, the Space Shuttle faced one hostile environment after another. One of those harsh environments—the hydrogen environment—existed within the shuttle itself. Liquid hydrogen was the fuel that powered the shuttle's complex, powerful, and reusable main engine. Hydrogen provided the high specific impulse—the bang per pound of fuel needed to perform the shuttle's heavy-lifting duties. Hydrogen, however, was also a potential threat to the very metal of the propulsion system that used it.

The diffusion of hydrogen atoms into a metal can make it more brittle and prone to cracking—a process called hydrogen embrittlement. This effect can reduce the toughness of carefully selected and prepared materials. A concern that exposure to hydrogen might encourage crack growth was present from the beginning of the Space Shuttle Program, but the rationale for using hydrogen was compelling.

The Challenge of the Hydrogen Environment

Hydrogen embrittlement posed more than a single engineering problem for the Space Shuttle. This was partly because hydrogen embrittlement can occur in three different ways. The most common mode occurs when hydrogen is absorbed by a material that is relatively unstressed, such as the components of the shuttle's main

engines before they experienced the extreme loads of liftoff and flight; this is called internal hydrogen embrittlement. Under the right conditions, internal hydrogen embrittlement has the potential to render materials too weak and brittle to survive high stresses applied later.

Alternatively, embrittlement can affect a material that is immersed in hydrogen while the material is being stressed and deformed. This phenomenon is called hydrogen environment embrittlement, which can occur in pressurized hydrogen storage vessels. These vessels are constantly stressed while in contact with hydrogen. Hydrogen environment embrittlement can potentially reduce ductility over time and enable cracking, or hydrogen may simply reduce the strength of a vessel until it is too weak to bear its own pressure.

Finally, hydrogen can react chemically with elements that are present in a metal, forming inclusions that can degrade the properties of that metal or even cause blisters on the metal's surface. This effect is called hydrogen reaction embrittlement. In the shuttle's main engine components, the reaction between hydrogen and the titanium alloys occurred to internally form brittle titanium hydrides, which was most likely to occur at locations where there were high tensile stresses in the part. Hydrogen reaction embrittlement can affect steels when hydrogen atoms combine with the carbon atoms dissolved in the metal. Hydrogen reaction embrittlement can also blister copper when hydrogen reacts with the internal oxygen in a solid copper piece, thereby forming steam blisters.

Insights on Hydrogen Environment Embrittlement

NASA studied the effects of hydrogen embrittlement in the 1960s. In the early 1970s, the scope of NASA-sponsored research broadened to include hydrogen environment embrittlement effects on fracture and fatigue. Engineers immersed specimens in hydrogen and performed a battery of tests. They applied repeated load cycles to specimens until they fatigued and broke apart; measured crack growth rates in cyclic loading and under a constant static load; and tested materials in high-heat and high-pressure hydrogen environments. Always, results were compared for each material to its performance in room-temperature air.

During the early years of the Space Shuttle Program, NASA and contractor engineers made a number of key discoveries regarding hydrogen environment embrittlement. First, cracks were shown to grow faster when loaded in a hydrogen environment. This finding would have significant implications for the shuttle design, as fracture assessments of the propulsion system would have to account for accelerated cracking. Second, scientists observed that hydrogen environment embrittlement could result in crack growth under a constant static load. This behavior was unusual for metals. Ductile materials such as metals tend to crack in alternating stress fields, not in fixed ones, unless a chemical or an environmental cause is present. Again, the design of the shuttle would have to account for this effect. Finally, hydrogen environment embrittlement was shown to have more severe effects at higher pressures. Intriguingly, degradation of tensile properties was found to be proportional to the square root of pressure.



The overall approach to hydrogen environment embrittlement research was straightforward. As a matter of common practice, NASA characterized the strength and fracture behavior of its alloys. To determine how these alloys would tolerate hydrogen, engineers simply adapted their tests to include a high-pressure hydrogen environment. After learning that high pressure exacerbates hydrogen environment embrittlement, they further adapted the tests to include a hydrogen pressure of 703 kg/cm² (10,000 psi). Later in the program, materials being considered for use in the main engine were tested at a reduced pressure of 492 kg/cm (7,000 psi) to be more consistent with operation conditions. The difference between room-temperature air material property data and these new results was a measurable effect of hydrogen environment embrittlement. Now that these effects could be quantified, the next step was to safeguard the shuttle.

Making Parts Resistant to Hydrogen Environment Embrittlement

One way to protect the main engines from hydrogen environment embrittlement was through materials selection. NASA chose naturally resistant materials when possible. There were, however, often a multitude of conflicting demands on these materials: they had to be lightweight, strong, tough, well suited for the manufacturing processes that shaped them, weldable, and able to bear significant temperature swings. The additional constraint of imperviousness to hydrogen environment embrittlement was not always realistic, so engineers

experimented with coatings and plating processes. The concept was to shield vulnerable metal from any contact with hydrogen. A thin layer of hydrogen environment embrittlement-resistant metal would form a barrier that separated at-risk material from hydrogen fuel.

Engineers concentrated their research on coatings that had low solubility and low-diffusion rates for hydrogen at room temperature. Testing had demonstrated that hydrogen environment embrittlement is worst at near-room temperature, so NASA selected coatings based on their effectiveness in that range. The most efficient barrier to hydrogen, engineers found, was gold plating; however, the cost of developing gold plating processes was a significant factor. Engineers observed that copper plating provided as much protection as gold, as long as a thicker and heavier layer was applied.

Protecting weld surfaces was often more challenging. The weld surfaces exposed to hydrogen fuel during flight were typically not accessible to plating after the weld was complete. Overcoming this problem required a more time-consuming and costly approach. Engineers developed weld overlays, processes in which hydrogen environment embrittlement-resistant filler metals were added during a final welding pass. These protective fillers sealed over the weld joints and provided the necessary barrier from hydrogen. NASA used overlays in combination with plating of accessible regions to prevent hydrogen environment embrittlement in engine welds.

These approaches—a combination of two or more hydrogen environment embrittlement prevention methods—were the practical solution for many of the embrittlement-vulnerable parts of the engines. For example, the most heavily used alloy in the engines was Inconel® 718, an alloy known to be affected by hydrogen environment embrittlement. Engineers identified an alternative heat treatment, different from the one typically used, which limited embrittlement. But this alone was insufficient. In the most critical locations, the alternative heat treatment was combined with copper plating and weld overlays.

A unique processing approach was also used to prevent embrittlement in the engine's main combustion chamber. This chamber was made with a highly conductive copper alloy. Its walls contained cooling channels that circulated cold liquid hydrogen and kept the chamber from melting in the extreme heat of combustion. But the hydrogen-filled channels became prone to hydrogen environment embrittlement. These liquid hydrogen channels were made by machining slots in the copper and then plated with nickel, which closed out the open slot and formed a coolant channel. The nickel plate cracked in the hydrogen environment and reduced the pressure capability of the channels. Engineers devised a two-part solution. First, they developed an alternative heat treatment to optimize nickel's performance in hydrogen. Next, they coated the nickel with a layer of copper to isolate it from the liquid hydrogen. This two-pronged strategy worked, and liquid hydrogen could be safely used as the combustion chamber coolant.



Addressing Internal Hydrogen Embrittlement

Whereas hydrogen environment embrittlement was of great concern at NASA in the 1960s, internal hydrogen embrittlement was largely dismissed even through the early years of the Space Shuttle Program. Internal hydrogen embrittlement had never been a significant problem for the types of materials used in spaceflight hardware. The superalloys and particular stainless steels selected by NASA were thought to be resistant to internal hydrogen embrittlement. Engineers thought the face-centered, cubic, close-packed crystal structure would leave too little room for hydrogen to permeate and diffuse.

Recall that internal hydrogen embrittlement occurs when hydrogen is absorbed before high operational stresses. Hydrogen enters into the metal and remains there, making it more brittle and likely to crack when extreme service loads are applied later. It is the accumulation of absorbed hydrogen, rather than the immediate exposure at the moment of high stress, that compromises an internal hydrogen embrittlement-affected material. When NASA initially designed the main engine, engineers accounted for hydrogen absorbed during manufacturing. Engineers, however, thought that the materials that were formed and processed without collecting a significant amount of hydrogen were not in danger of absorbing considerable amounts later.

This notion about internal hydrogen embrittlement was challenged during the preparation of an engine failure analysis document in 1988. The engine

was repeatedly exposed to hydrogen in flight and after flight, at high temperatures and extreme pressure. The report suggested that in these exceptional heat and pressure conditions some engine materials might, in fact, gather small amounts of hydrogen with each flight. Gradually, over time, these materials could accumulate enough hydrogen to undermine ductility.

Engineers developed a special test regimen to screen materials for high-temperature, high-pressure hydrogen accumulation. Test specimens were “charged” with hydrogen at 649°C (1,200°F) and 351.6 kg/cm² (5,000 psi). They were then quickly cooled and tested for strength and ductility under normal conditions. Surprisingly, embrittlement by internal hydrogen embrittlement was observed to be as severe as by hydrogen environment embrittlement. As a subsequent string of fatigue tests confirmed this comparison, NASA had to reevaluate its approach to preventing hydrogen embrittlement. The agency’s focus on hydrogen environment embrittlement had been a near-total focus. Now, a new awareness of internal hydrogen embrittlement would drive a reexamination.

Fortunately, the process for calculating design properties from test data had been conservative. The margins of safety were wide enough to bound the combined effects of internal hydrogen embrittlement and hydrogen environment embrittlement. The wealth of experience gained in studying hydrogen environment embrittlement and mitigating its effects also worked in NASA’s favor. Some of the same methodologies could now be applied to

internal hydrogen embrittlement. For instance, protective plating would operate on the same principle—the creation of a barrier between hydrogen and a vulnerable alloy—whether hydrogen environment embrittlement or internal hydrogen embrittlement was the chief worry. Continued testing of “charged” specimens would allow quantification of internal hydrogen embrittlement damage, just as hydrogen immersion testing had enabled measurement of hydrogen environment embrittlement effects.


Taking strategies generated to avoid hydrogen environment embrittlement and refitting them to prevent internal hydrogen embrittlement, however, often required additional analysis. For example, from the beginning of the Space Shuttle Program NASA used coatings to separate at-risk metals from hydrogen. The agency intentionally chose these coatings for their performance at near-room temperature, when hydrogen environment embrittlement is most aggressive. Tests showed the coatings were less effective in the high heat that promotes internal hydrogen embrittlement. New research and experimentation was required to prove that these protective coatings were adequate—that, although they didn’t completely prevent the absorption of hydrogen when temperatures and pressures were extreme, they did reduce it to safe levels.

Special Cases: High-Pressure Fuel Turbopump Housing

NASA encountered a unique hydrogen embrittlement issue during development testing of the main engine high-pressure fuel turbopump.




High-Pressure Fuel and Oxidizer Turbopump Turbine Blade Cracks




After observing cracks on polycrystalline turbine blades, NASA redesigned the blades as single-crystal parts. When tested in hydrogen, cracks were detected. Scientists used a Brazilian disc test to create the tensile and shear stresses that had caused growth. NASA resolved cracking in the airfoil with changes that eliminated stress concentrations and smoothed the flow of molten metal during casting. To assess cracking at damper contacts, scientists extracted test specimens from single crystal bars, machined contact pins from the damper material, and loaded two specimens. This contact fixture was supported in a test rig that allowed the temperature, loads, and load cycle rate to be varied. Specimens were pre-charged with hydrogen, tested at elevated temperatures, and cycled at high frequency to actual operating conditions.



First Stage Blade 42 Trailing Edge Root



Schematic of Test Rig





A leak developed during the test; this leak was traced to cracks in the mounting flange of the turbopump's housing. The housing was made from embrittlement-prone nickel-chromium alloy Inconel® 718, and the cracks were found to originate in small regions of highly concentrated stress. So, engineers changed the material to a more-hydrogen-tolerant alloy, Inconel® 100, and they redesigned the housing to reduce stress concentrations. This initially appeared to solve the problem. Then, cracks were discovered in other parts of the housing. Structural and thermal analysis could not explain this cracking. The locations and size of the cracks did not fit with existing fatigue and crack-growth data.

To resolve this inconsistency, engineers considered the service conditions of the housing. The operating environment of the cracked regions was a mixture of high-pressure hydrogen and steam at 149°C to 260°C (300°F to 500°F). Generally, hydrogen environment embrittlement occurs near room temperature and would not be a significant concern at that level of heat; however, because of the unexplained cracking, a decision was made to test Inconel® 100 at elevated temperatures in hydrogen and hydrogen mixed with steam. Again, the results were unexpected. Engineers observed a pronounced reduction in strength and ductility in these environments at elevated temperatures. Crack growth occurred at highly accelerated rates—as high as two orders of magnitude above room-temperature air when the crack was heavily loaded to 30 ksi $\sqrt{\text{in}}$ (33 MPa $\sqrt{\text{m}}$) and held for normal engine operating time. Moreover, crack growth was driven by both the number of load cycles and the duration of each load cycle. Crack growth is typically sensitive to the number and magnitude of load cycles but not to the length of time for each cycle.

Clearly, the combination of the hydrogen and steam mixture and the uncommonly high stress concentrations was promoting hydrogen environment embrittlement in Inconel® 100 at high temperatures. Resolving this issue required three modifications. First, detailed changes to the shape of the housing were made, further reducing stress concentrations. Second, gold plating was added to shield the Inconel® 100 from the hot hydrogen and steam mixture. Finally, a manufacturing process called “shot peening” was used to fortify the surface of the housing against tensile stresses by impacting it with shot, determined to be promoting fracture, and therefore eliminated.

Summary

The material characterization done in the design phase of the main engine, and the subsequent anomaly resolution during its development phase, expanded both the material properties database and the understanding of hydrogen embrittlement. The range of hydrogen embrittlement data has been broadened from essentially encompassing only steels to now including superalloys. It was also extended from including primarily tensile properties to including extensive low-cycle fatigue and fracture-mechanics testing in conditions favorable to internal hydrogen embrittlement or hydrogen environment embrittlement. The resultant material properties database, now approaching 50 years of maturity, is valuable not only because these materials are still being used, but also because it serves as a foundation for predicting how other materials will perform under similar conditions—and in the space programs of the future.

Space Environment: It's More Than a Vacuum

We know that materials behave differently in different environments on Earth. For example, aluminum does not change on a pantry shelf for years yet rapidly corrodes or degrades in salt water.

One would think that such material degradation effects would be eliminated by going to the near-perfect vacuum of space in low-Earth orbit. In fact, many of these effects are eliminated. However, Orbiter systems produced gas, particles, and light when engines, overboard dumps, and other systems operated, thereby creating an induced environment in the immediate vicinity of the spacecraft. In addition, movement of the shuttle through the tenuous upper reaches of Earth's atmosphere (low-Earth orbit) at orbital velocity produced additional contributions to the induced environment in the form of spacecraft glow and atomic oxygen effects on certain materials. The interactions of spacecraft materials with space environment factors like solar ultraviolet (UV) light, atomic oxygen, ionizing radiation, and extremes of temperature can actually be detrimental to the life of materials used in spacecraft systems.

For the Orbiter to perform certain functions and serve as a platform for scientific measurements, the effects of natural and Orbiter-induced environments had to be evaluated and controlled. Payload sensitivities to these environmental effects varied, depending on payload characteristics. Earth-based observatories and other instruments are affected by the Earth's atmosphere in terms of producing unwanted light background and other contamination effects. Therefore, NASA developed



essential analytical tools for environment prediction as well as measurement systems for environment definition and performance verification, thus enabling a greater understanding of natural and induced environment effects for space exploration.

Induced Environment Characterization

NASA developed mathematical models to assess and predict the induced environment in the Orbiter cargo bay during the design and development phase of the Space Shuttle Program. Models contained the vehicle geometry, vehicle flight attitude, gas and vapor emission source characteristics, and used low-pressure gas transport physics to calculate local gas densities, column densities (number of molecular species seen along a line of sight), as well as contaminant deposition effects on functional surfaces. Gas transport calculations were based on low-pressure molecular flow physics and included scattering from Orbiter surfaces and the natural low-Earth orbit environment.

The Induced Environment Contamination Monitor measured the induced environment on three missions—Space Transportation System (STS)-2 (1981), STS-3 (1982), and STS-4 (1982)—and was capable of being moved using the Shuttle Robotic Arm to various locations for specific measurements. Most measurements were made during the on-orbit phase. This measurement package was flown on the three missions to assess shuttle system performance. Instruments included a humidity monitor, an air sampler for gas collection and analysis after return, a cascade impactor for particulate measurement, passive samples for optical degradation of



The Atlantic Ocean southeast of the Bahamas is in the background as Columbia's Shuttle Robotic Arm and end effector grasp a multi-instrument monitor for detecting contaminants. The experiment, called the Induced Environment Contaminant Monitor, was flown on STS-4 (1982). The tail of the Orbiter can be seen below.

surfaces, quartz-crystal microbalances for deposited mass measurement, a camera/photometer pair for particle measurement in the field of view, and a mass spectrometer. Additional flight measurements made on STS-52 (1992) and many payloads provided more data.

Before the induced environment measurements could be properly interpreted, several on-orbit operational aspects needed to be understood. Because of the size of the vehicle and its payloads, desorption of adsorbed gases such as water, oxygen, and nitrogen (adsorbed on Earth) took a fairly long time, the induced environment on the first day of a mission was affected more than on

subsequent days. Shuttle flight attitude requirements could affect the cargo bay gaseous environment via solar heating effects as well as the gases produced by engine firings. These gases could reach the payload bay by direct or scattered flow. Frequently, specific payload or shuttle system attitude or thermal control requirements conflicted with the quiescent induced environment required by some payloads.

With the above operational characteristics, data collected with the monitor and subsequent shuttle operations showed that, in general, the measured data either met or were close to the requirements of sensitive payloads during quiescent periods. A large qualification to this statement



had to be made based on a new understanding of the interaction of the natural environment with vehicle surfaces. This interaction resulted in significantly more light emissions and material surface effects than originally expected. Data also identified an additional problem of recontact of particles released from the shuttle during water dumps with surfaces in the payload bay. The induced environment control program instituted for the Space Shuttle Program marked a giant step from the control of small free-flying instrument packages to the control of a large and complex space vehicle with a mixed complement of payloads. This approach helped develop a system with good performance, defined the vehicle associated environment, and facilitated effective communication between the program and users.

The induced environment program also showed that some attached payloads were not compatible with the shuttle system and its associated

payloads because of the release of water over long periods of time. Other contamination-sensitive payloads such as Hubble Space Telescope, however, were not only successfully delivered to space but were also repaired in the payload bay.

Unique Features Made It Possible


The Orbiter was the first crewed vehicle to provide protection of instrumentation and sensitive surfaces in the payload bay during ascent and re-entry and allow exposure to the low-Earth orbit environment. Effects were observed without being modified by flight heating or gross contamination. Also, as part of the induced environment control program, the entire payload bay was examined immediately on return. Because of these unique aspects, NASA was able to discover and quantify unexpected interactions between the environment of low-Earth and the vehicle.

Discovery of Effects of Oxygen Atoms

After STS-1 (1981) returned to Earth, researchers visually examined the material surfaces in the payload bay for signs of contamination effects. Most surfaces appeared pristine, except for the exterior of the television camera thermal blankets and some painted surfaces. The outside surface of the blankets consisted of an organic (polyimide) film that, before flight, appeared gold colored and had a glossy finish. After flight, most films were altered to a yellow color and no longer had a glossy finish but, rather, appeared carpet-like under high magnification. Only the surfaces of organic materials were affected; bulk properties remained unchanged.

Patterns on modified surfaces indicated directional effects and, surprisingly, the flight-exposed surfaces were found to have receded rather than having deposited contaminants. The patterns on the surfaces were related to the

Atomic Oxygen Effects on Polymers and Plastics in low-Earth Orbit as Seen With the Scanning Electron Microscope; STS-46 (1992)



- Scanning electron microscope image of a typical Kapton® polyimide plastic sheet. The various specs and bumps are from the inorganic filler used in plastic sheet manufacture.
- Scanning electron microscope image of a typical Kapton® polyimide plastic sheet after exposure to surface bombardment by atomic oxygen in low-Earth orbit. The rough surface is typical of atomic oxygen attack on plastics in low-Earth orbit and is the result of the strong dependence of chemical reaction on atom-surface collision energy. Note how some of the inorganic filler particles are standing on pedestals because they protect the underlying plastic from atomic oxygen attack.
- Scanning electron microscope image of a microelectron fabrication etching target also flown on STS-46 and exposed to low-Earth orbit atomic oxygen. The highly directional attack of low-Earth orbit atomic oxygen produced a clean, high-resolution removal of the unprotected plastic around the pattern of protective inorganic surface coatings. High-speed neutral atomic oxygen beams in ground-based production facilities may be a useful adjunct to microelectronic production as described in US Patent 5,271,800.



vehicle velocity vector. When combining these data with the atmospheric composition and densities, the material surface recession was caused by the high-velocity collision of oxygen atoms with forward-facing Orbiter surfaces leading to surface degradation by oxidation reactions. Oxygen atoms are a major constituent of the natural low-Earth orbit environment through which the shuttle flew at an orbital velocity of nearly 8 km/sec (17,895 mph). The collision energy of oxygen atoms striking forward-facing shuttle surfaces in low-Earth orbit was extremely high—on the order of 5 electron volts (eV)—100 times greater than the energy of atoms in typical low-pressure laboratory oxygen atom generators. The high collision energy of oxygen atoms in low-Earth orbit plays an important role in surface reactivity and surface recession rates.

Material recession rates are determined by normalizing the change in sample mass to the number of oxygen atoms reaching the surface over the exposure time (atoms/cm², fluence). Atom density is obtained from the standard atmospheric density models used by NASA and the Department of Defense. Since oxygen atoms travel much slower than the Orbiter, they impacted the surfaces in question only when facing toward the vehicle velocity vector and had to be integrated over time and vehicle orientation. STS-1 recession data were approximate because they had to be integrated over changing vehicle attitude; had limited atom flux, uncontrolled surface temperatures and solar UV exposure; and predicted atom densities. Recession rates determined from material samples exposed during the STS-5 (1982) mission and Induced Environmental Contamination Monitor

flights had the same limitations but supported the STS-1 data. Extrapolation of these preliminary recession data to longer-term missions showed the potential for significant performance degradation of critical hardware, so specific flight experiments were carried out to quantify the recession characteristics and rates for materials of interest.

On-orbit Materials Behavior

Fifteen organizations participated in a flight experiment on STS-8 (1983) to understand materials behavior in the low-Earth orbit environment. The objective was to control some of the parameters to obtain more-accurate recession rates. The mission had a dedicated exposure to direct atom impact (payload bay pointing in the velocity direction) of 41.7 hours at an altitude of 225 km (121 nautical miles) resulting in the largest fluence of the early missions (3.5×10^{20} atoms/cm²). Temperature control at two set points was provided as well as instruments to control UV and exposure to electrically charged ionospheric plasma species.

The STS-8 experiment provided significant insight into low-Earth orbit environment interactions with materials. Researchers established quantitative reaction rates for more than 50 materials, and were in the range of $2\text{--}3 \times 10^{-24}$ cm³/atom for hydrocarbon-based materials. Perfluorinated organic materials were basically nonreactive and silicone-based materials stopped reacting after formation of a protective silicon oxide surface coating. Material reaction rates, as a first approximation, were found to be independent of temperature, material morphology, and exposure to solar radiation or electrically charged ionospheric species.

Researchers also evaluated coatings that could be used to protect surfaces from interaction with the environment.

Reaction rates were based on atomic oxygen densities determined from long-term atmospheric density models, potentially introducing errors in short-term experiment data. In addition, researchers obtained very little insight into the reaction mechanism(s).

An additional flight experiment—Evaluation of Oxygen Interaction with Materials III—addressing both of these questions was flown on STS-46 (1992). The primary objective was to produce benchmark atomic oxygen reactivity data by measuring the atom flux during material surface exposure. Secondary experiment objectives included: characterizing the induced environment near several surfaces; acquiring basic chemistry data related to reaction mechanism; determining the effects of temperature, mechanical stress, atom fluence, and solar UV radiation on material reactivity; and characterizing the induced and contamination environments in the shuttle payload bay. This experiment was a team effort involving NASA centers, US Air Force, NASA Space Station Freedom team, Aerospace Corporation, University of Alabama in Huntsville, National Space Agency of Japan, European Space Agency, and the Canadian Space Agency.


STS-46 provided an opportunity to make density measurements at several altitudes: 427, 296, and 230 km (231, 160, and 124 nautical miles). However, the vehicle flew for 42 hours at 230 km (124 nautical miles) with the payload bay surfaces pointed into the velocity vector during the main portion of the mission to obtain high fluence. The mass spectrometer provided by the



US Air Force was the key component of the experiment and was capable of sampling both the direct atomic oxygen flux as well as the local neutral environment created by interaction of atomic oxygen with surfaces placed in a carousel. Five carousel sections were each coated with a different material to determine the material effects on released gases. Material samples trays, which provided temperature control plus instruments to control other exposure conditions, were placed on each side of the mass spectrometer/carousel.

NASA achieved all of the Evaluation of Oxygen Interaction with Materials III objectives during STS-46. A well-characterized, short-term, high-fluence atomic oxygen exposure was provided for a large number of materials, many of which had never been exposed to a known low-Earth orbit atomic oxygen environment. The data provided a benchmark reaction rate database, which has been used by the International Space Station, Hubble, and others to select materials and coatings to ensure long-term durability.

Reaction rate data for many of the materials from earlier experiments were confirmed, as was the generally weak dependence of these reaction rates on temperature, solar UV exposure, oxygen atom flux, and exposure to charged ionospheric species. The role of surface collision energy on oxygen atom reactivity was quantified by comparing flight reaction rates of key Evaluation of Oxygen Interaction with Materials III experiment materials with reactivity measurements made in well-characterized laboratory oxygen atom systems with lower surface collision energies. This evaluation also provided an important benchmark point for understanding the role of



Evaluation of Oxygen Interaction with Materials III flight experiment in the Orbiter payload bay of STS-46 (1992). Material exposure samples are located on both sides of the mass spectrometer gas evolution measurement assembly in the center.

solar extreme UV radiation damage in increasing the generally low surface reactivity of perfluorinated organic materials. The mass spectrometer/carousel experiment produced over 46,000 mass spectra providing detailed characterization of both the natural and the induced environment. The mass spectrometer database provided a valuable resource for the verification of various models of rarified gas and ionospheric plasma flow around spacecraft.

Intelsat Satellite

Knowledge gained from atomic oxygen reactivity studies played a key role in the STS-49 (1992) rescue of the communications satellite

Intelsat 603 that was used to maintain communications from a geosynchronous orbit. Failure of the Titan-3 upper stage left Intelsat 603 marooned in an unacceptable low-Earth orbit and subject to the effects of atomic oxygen degradation of its solar panels, which could have rendered the satellite useless. NASA quickly advised the International Telecommunications Satellite Organization (Intelsat) Consortium of the atomic oxygen risk to Intelsat 603, leading to the decision to place the satellite in a configuration that was expected to minimize atomic oxygen damage to the silver interconnects on the solar panels. This was accomplished by raising the satellite altitude and changing its flight attitude so that atomic oxygen fluence was minimized.



The Intelsat Solar Array Coupon flight experiment shown mounted on the Shuttle Robotic Arm lower arm boom and exposed to space environment conditions during STS-41 (1990).

To provide facts needed for a final decision about a rescue flight, NASA designed and executed the Intelsat Solar Array Coupon flight experiment on STS-41 (1990). The experiment results, in combination with ground-based testing, supported the decision to conduct the STS-49 satellite rescue mission. On this mission, Intelsat 603 was captured and equipped with a solid re-boost motor to carry it to successful geosynchronous orbit.


STS-62 (1994) orbits Earth during a "night" pass, documenting the glow phenomenon surrounding the vertical stabilizer and the Orbital Maneuvering System pods of the spacecraft.

NASA Discovers Light Emissions

On the early shuttle flights, NASA observed another effect caused by the interaction between spacecraft surfaces and the low-Earth orbit environment. Photographs obtained by using intensified cameras and conducted from the Orbiter cabin windows showed light emissions (glow) from the Orbiter surfaces when in forward-facing conditions.

The shuttle provided an excellent opportunity to further study this phenomenon. On STS-41D (1984), astronauts photographed various material samples using a special glow spectrometer to obtain additional data and determine if the glow was dependent on surface composition. These measurements, along with the material recession effects and data obtained on subsequent flights, led to a definition of the glow mechanism.

Spacecraft glow is caused by the interaction of high-velocity oxygen atoms with nitrous oxide absorbed on the surfaces, which produces nitrogen dioxide in an electronically excited state. The excited nitrogen dioxide is released from the surfaces and emits light as it moves away and decays from its excited state. Some nitrous oxide on the surface and some of the released nitrogen dioxide result from the natural environment. The light emission occurs on any spacecraft operating in low-Earth orbit; however, the glow could be enhanced by operation of the shuttle attitude control engines, which produced nitrous oxide and nitrogen dioxide as reaction products. These findings led to a better understanding of the behavior of spacecraft operating in low-Earth orbit and improved accuracy of instrument measurements.





Chemical Fingerprinting

Comprehensive Electronic System for Greater Flight Safety

A critical concern for all complex manufacturing operations is that contaminants and material changes over time can creep into the production environment and threaten product quality. This was the challenge for the solid rocket motors, which were in production for 30 years.

It is possible that vendor-supplied raw materials appear to meet specifications

from lot to lot and that supplier process changes or even contaminated material can appear to be “in spec” but actually contain subtle, critical differences. This situation has the potential to cause significant problems with hardware performance.

NASA needed a system to readily detect those subtle yet potentially detrimental material variances to ensure the predictability of material properties and the reliability of shuttle reusable solid rocket motors. The envisioned solution was to pioneer consistent and repeatable analytical methods tailored to specific, critical materials that would yield accurate assessments of material

integrity over time. Central to the solution was both a foolproof analysis process and an electronic data repository for benchmarking and monitoring.


A Chemical “Fingerprint”

Just as fingerprints are a precise method to confirm an individual’s identity, the solid rocket motor project employed chemical “fingerprints” to verify the quality of an incoming raw material. These fingerprints comprised a detailed spectrum of a given material’s chemical signature, which could be captured digitally and verified using a combination of sophisticated laboratory equipment and custom analytical methods.

The challenge was to accurately establish a baseline chemical fingerprint of each material and develop reproducible analytical test methods to monitor lot-to-lot material variability. A further objective was to gain a greater understanding of critical reusable solid rocket motor materials, such as insulation and liner ingredients, many of which were the same materials used since the Space Shuttle Program’s inception. New analytical techniques such as the atomic force microscope were used to assess materials at fundamental chemical, molecular, and mechanical levels. These new techniques provided the high level of detail sought. Because of unique attributes inherent in each material, a one-size-fits-all analysis method was not feasible.

To facilitate documentation and data sharing, the project team envisioned a comprehensive electronic database to provide ready access to all relevant data. The targeted level of background detail included everything from where and how a material was properly used to details of chemical composition.


Environmental Assurance



During the Space Shuttle Program’s operation, issues arose regarding the use of substances that did not meet emerging environmental regulations and current industry standards. NASA worked to develop chemicals, technologies, and processes that met regulatory requirements, and the agency strove to identify, qualify, and replace materials that were becoming obsolete as a result of environmental issues. The stringent demands of human spaceflight required extensive testing and qualification of these replacement materials.



Tools for Materials Evaluation Atomic Force Microscope Images of Metal Surface



The atomic force microscope affords a visual evaluation of surface preparation processes to improve understanding of their effects on bonding. The top panel represents topography of a grit blast surface for comparison to a highly polished one. The atomic force microscope uses an extremely fine probe to measure minute interactions with surface features even down to an atomic scale. The maps at left are scaled from black at the bottom of valleys to white at the tops of peaks within the scanned area. The 3-D projections at right are on a common height scale. The grit blast surface clearly offers greatly increased surface area and mechanical interlocking for enhanced bonding. Beyond simple topography, the probe interactions with atomic forces can also measure and map properties such as microscopic hardness or elastic modulus on various particles and/or phase transitions in a composite material, which in turn can be correlated with chemical and physical properties.

© ATK. All rights reserved.

The ideal system would enable a qualified chemist to immediately examine original chemical analysis data for the subtle yet significant differences between the latest lot of material and previous good or bad samples.

To develop such a system, commercially available hardware and software were used to the greatest extent possible. Since an electronic framework to tie the data together did not exist, one was designed in-house.

The Fingerprinting Process

The chemical fingerprinting program, which began in 1998 with a prioritized list of 14 critical materials, employed a team approach to quantify and document each material. The interdisciplinary team included design engineering, materials and processes engineering, procurement quality engineering, and analytical chemistry. Each discipline group proposed test plans that included the types of testing to be developed. Following approval,

researchers acquired test samples (usually three to five lots of materials) and developed reliable test methods. Because of the unique nature of each material, test methods were tailored to each of the 14 materials.

A “material” site in the project database was designed to ensure all data were properly logged and critical reports were written and filed. Once the team agreed sufficient data had been generated, a formal report was drafted and test methods were selected to develop new standard acceptance procedures that would ultimately be used by quality control technicians to certify vendor materials.

The framework developed to package the wide-ranging data was termed the Fingerprinting Viewer. Program data were presented through a series of cascading menu pages, each with increasing levels of detail.

The Outcomes

Beyond meeting the primary program objectives, a number of resulting benefits were noted. First, through increased data sharing, employees communicated more effectively, both internally and with subtier suppliers. The powerful analytical methods employed also added to the suppliers’ materials knowledge base. Subtle materials changes that possibly resulted from process drift or changes at subtier suppliers were detectable. Eight subtier suppliers subsequently implemented their own in-house chemical fingerprinting programs to improve product consistency, recertify material after production changes, or even help develop key steps in the manufacturing process to ensure repeatable quality levels.

Additionally, engineers could now accurately establish shelf-life extensions and storage requirements



This high-performance liquid chromatography/mass spectrometry is employed to document minute details of a material's chemical and molecular composition. Through the chemical fingerprinting system, seemingly minuscule discrepancies raise red flags that trigger investigations and preclude defective materials from reaching the production floor. Dr. Ping Li shown here at ATK in Utah.

for stockpiled materials. The ability to store greater amounts of materials over longer periods of time was valuable in cases where new materials needed to be certified to replace existing materials that had become obsolete.

Finally, investigators were able to solve production issues with greater efficiency. Comprehensive database features, including standardized test methods and the extensive online reference database, provided resources needed to resolve production issues in a matter of days or even hours—issues that otherwise would have required major investigations. In some cases, fingerprinting was also used to indicate that a suspect material was actually within required specifications. These materials may have been rejected in previous cases but, by using the fingerprinting database to assess the

material, the team could look deeper to find the true root cause and implement proper corrective actions.

From Fingerprints to Flight Safety

The overarching value of the chemical fingerprinting program was that it provided greater assurance of the safety and reliability of critical shuttle flight hardware. The fundamental understanding of critical reusable solid rocket motor materials and improved communications with vendors reduced the occurrence of raw materials issues. NASA will implement chemical fingerprinting methods into the acceptance testing of raw materials used in future human space exploration endeavors. The full benefits of the program will continue to be realized in years to come.

Unprecedented Accomplishments in the Use of Aluminum-Lithium Alloy

NASA was the first to use welded aluminum-lithium alloy Al 2195 at cryogenic temperatures, incorporating it into the External Tank under circumstances that demanded innovation.

From the beginning of the Space Shuttle Program's launch phase, NASA sought to reduce the weight of the original tank, thereby increasing payload capacity. Since the tank was carried nearly to orbit, close to 100% of the weight trimmed could be applied to the payload. NASA succeeded in implementing numerous weight-saving measures, but the biggest challenge was to incorporate a lightweight aluminum alloy—aluminum-lithium Al 2195—into the tank structure. This alloy had never been used in welded cryogenic environments prior to NASA's initiative. Several challenges needed to be overcome, including manufacturing the aluminum-lithium tank components, welding the alloy, and repairing the welds. NASA and the External Tank prime contractor broke new ground in the use of aluminum-lithium to produce the "super lightweight tank."

The original tank weighed 34.500 metric tons (76,000 pounds) dry. By the sixth shuttle mission, the tank's weight had been reduced to 29.900 metric tons (66,000 pounds). This configuration was referred to as the "lightweight tank."

The real challenge, however, was still to come. In 1993, the International Space Station Program decided to change the station's orbital inclination



to 57 degrees (a “steeper” launch inclination), allowing Russian vehicles to fly directly to the station. That change cost the shuttle 6,123 kg (13,500 pounds) of payload capacity. The External Tank project office proposed to reduce the dry weight of the tank by 3,402 kg (7,500 pounds).

The Space Shuttle Program sought to incorporate lightweight aluminum-lithium Al 2195 into the majority of the tank structure, replacing the original aluminum-copper alloy Al 2219; however, NASA first needed to establish requirements for manufacturing, welding, and repairing aluminum-lithium weld defects.

NASA started the super lightweight tank program in 1994. During the early phase, advice was sought from welding experts throughout the United States and the United Kingdom. The consensus: it was virtually impossible to perform repairs on welded aluminum-lithium.

The aluminum-lithium base metal also presented challenges. Lockheed Martin worked with Reynolds Aluminum to produce the aluminum-lithium base metal. One early problem was related to aluminum-lithium material’s fracture toughness—a measure of the ability of material with a defect to carry loads. Although material was screened, flight hardware requirements dictated that structures must have the ability to function in the event a defect was missed by the screening process. The specific difficulty with the aluminum-lithium was that the cryogenic fracture toughness of the material showed little improvement over the room-temperature fracture toughness.

Since the two propellant tanks were proof tested at room temperature and flown cryogenically, this fracture toughness ratio was a crucial factor.

A simulated service test requirement was imposed as part of lot acceptance for all aluminum-lithium material used on the tank. The test consisted of applying room temperature and cryogenic load cycles to a cracked sample to evaluate the ability of the material to meet the fracture toughness requirements. Failure resulted in the plate being remelted and reprocessed.

Implementation of simulated service testing as a lot acceptance requirement was unique to the aluminum-lithium material. Testing consisted of cropping two specimens from the end of each plate. Electrical discharge machining (a process that removes metal by discharging a spark between the tool and the test sample) was used to introduce a fine groove in each sample. The samples were then cyclically loaded at low stresses to generate a sharp fatigue crack that simulated a defect in the material.

The first sample was stressed to failure; the second sample was stressed to near failure and then subjected to cyclic loading representative of load cycles the tank would see on the launch pad during tanking and during flight.

In the second sample, initial loading was conducted at room temperature. This simulated the proof test done on the tank. Next, the sample was stressed 13 times (maximum tanking requirement) to the level expected during loading of propellants at cryogenic temperatures and, finally, stressed to maximum expected flight

stress at cryogenic temperature.

This cycle was repeated three more times to meet a four-mission-life program requirement with the exception that, on the fourth cycle, the sample was stressed to failure and had to exceed a predetermined percent of the flight stress. Given the size of the barrel plates for the liquid hydrogen and liquid oxygen tanks, only one barrel plate could be made from each lot of material. As a result, this process was adopted for every tank barrel plate—32 in each liquid hydrogen tank and four in each liquid oxygen tank—and implemented for the life of the program.

Another challenge was related to the aluminum-lithium weld repair process on compound curvature parts. The effect of weld shrinkage in the repairs caused a flat spot, or even a reverse curvature, in the vicinity of the repairs and contributed to significant levels of residual stress in the repair. Multiple weld repairs, in proximity, showed the propensity for severe cracking. After examination of the repaired area, it was found that welding aluminum-lithium resulted in a zone of brittle material surrounding the weld. Repeated repairs caused this zone to grow until the residual stress from the weld shrinkage exceeded the strength of the weld repair, causing it to crack.

The technique developed to repair these cracks was awarded a US Patent. The repair approach consisted of alternating front-side and back-side grinds as needed to remove damaged microstructure. It was also found that aluminum-lithium could not tolerate as much heating as the previous aluminum-copper alloy. This required increased torch speeds and decreased



The use of aluminum-lithium Al 2195 in manufacturing major External Tank components, such as the liquid hydrogen tank structure shown above, allowed NASA to reduce the overall weight of the External Tank by 3,402 kg (7,500 pounds). The liquid hydrogen tank measured 8.4 m (27.5 ft) in diameter and 29.4 m (96.6 ft) in length. Photo taken at NASA's Michoud Assembly Facility in New Orleans, Louisiana.

fill volumes to limit the heat to which the aluminum-lithium was subjected.

Additional challenges in implementing effective weld repairs caused NASA to reevaluate the criteria for measuring the strength of the welds. In general, weld repair strengths can be evaluated by excising a section of the repaired material and performing a tensile test. The strength behavior of the repaired material is compared to the strength behavior of the original weld material. In the case of the aluminum-copper alloy Al 2219, the strengths were

comparable; however, in the case of the aluminum-lithium alloy repair, the strengths were lower.

Past experience and conventional thinking was that in the real hardware, where the repair is embedded in a long initial weld, the repaired weld will yield and the load will be redistributed to the original weld, resulting in higher capability. To demonstrate this assumption, a tensile test was conducted on a 43-cm- (17-in.)-wide aluminum-lithium panel that was fabricated by welding two

aluminum-lithium panels together and simulating a weld repair in the center of the original weld. The panel was then loaded to failure. The test that was supposed to indicate better strength behavior than the excised repair material actually failed at a lower stress level.

To understand this condition, an extensive test program was initiated to evaluate the behavior of repairs on a number of aluminum-copper alloy (Al 2219) and aluminum-lithium alloy (Al 2195) panels.




Orbiter Payload Bay Door

One of the largest aerospace composite applications of its time.

With any space vehicle, minimum weight is of critical importance. Initial trade studies indicated that using a graphite/epoxy structure in place of the baselined aluminum structure provided significant weight savings of about 408 kg (900 pounds [4,000 newtons]), given the large size and excellent thermal-structural stability. Two graphite/epoxy composite materials and four structural concepts—full-depth honeycomb sandwich, frame-stiffened thin sandwich, stiffened skin with frames and stringers, and stiffened skin with frames only—were considered for weight savings and manufacturing producibility efficiency. These studies resulted in the selection of the frame-stiffened thin sandwich configuration, and component tests of small specimens finalized the graphite fiber layup, matrix material, and honeycomb materials. Graphite/epoxy properties at elevated temperatures are dependent on moisture content and were taken into account in developing mechanical property design allowables. Additionally, NASA tracked the moisture content through all phases of flight to predict the appropriate properties during re-entry when the payload bay doors encountered maximum temperatures of 177°C (350°F).

Payload bay doors were manufactured in 4.57-m (15-ft) sections, resulting in two 3 x 18.3 m (10 x 60 ft) doors. The panel face sheets consisted of a ± 45 -degree fabric ply imbedded between two 0-degree tape plies directed normal to the frames and were pre-cured prior to bonding to the Nomex® honeycomb core. A lightweight-aluminum wire mesh bonded to the outside of face sheets provided lightning-strike protection. Frames consisted primarily of fabric plies with the interspersions of 0-degree plies dictated by strength and/or



stiffness. Mechanical fasteners were used for connection of major subassemblies as well as final assembly of the doors.

All five Orbiter vehicles used graphite/epoxy doors, one of the largest aerospace composite applications at the time, and performance was excellent throughout all flights. Not only was the expected weight saving achieved and thermal-structural stability was acceptable, NASA later discovered that the graphite/epoxy material showed an advantage in ease of repair. Ground handling damage occurred on one section of a door, resulting in penetration of the outer skin of the honeycomb core. The door damage was repaired in 2 weeks, thereby avoiding significant schedule delay.



Test panels were covered with a photo-stress coating that, under polarized light, revealed the strain pattern in the weld repair. The Al 2219 panel behaved as expected: the repair yielded, the loads redistributed, and the panel pulled well over the minimum allowable value. In aluminum-lithium panels, however, the strains remained concentrated in the repair. Instead of the 221 MPa (32,000 pounds/in²) failure stress obtained in the initial welds, the welds were failing around 172 MPa (18,000 pounds/in²). These lower failure stress values were problematic due to a number of flight parts that had already been sized and machined for the higher 221 MPa (32,000 pounds/in²) value.

Based on this testing, it was determined that weld shrinkage associated with the repair resulted in residual stresses in the joint, reducing the joint capability. To improve weld repair strengths, engineers developed an approach to planish (lightly hammer) the weld bead, forcing it back into the joint and spreading the joint to redistribute and reduce the residual stresses due to shrinkage. This required scribing and measuring the joint before every repair, making the repair, and then planishing the bead to restore the weld to its previous dimensions. Wide panel test results and photo-stress evaluation of planished repairs revealed that the newly devised repair procedure was effective at restoring repair strengths to acceptable levels.

Testing also revealed that planishing of weld beads is hard to control precisely, resulting in the process frequently forming other cracks, thus leading to additional weld repairs. Because of the

difficulty in making and planishing multiple repairs, a verification ground rule was established that every “first repair of its kind” had to be replicated on three wide tensile panels, which were then tested either at room temperature or in a cryogenic environment, depending on the in-flight service condition expected for that part of the tank.

All these measures combined accomplished the first-ever use of welded aluminum-lithium at cryogenic temperatures, meeting the strict demands of human spaceflight. The super lightweight tank incorporated 20 aluminum-lithium ogive gores (the curved surfaces at the forward end of the liquid oxygen tank), four liquid oxygen barrel panels, 32 liquid hydrogen barrel panels, 12 liquid oxygen tank aft dome gores, 12 liquid hydrogen tank forward dome gores, and 11 liquid hydrogen aft dome gores.

Through this complex and innovative program, NASA reduced the 29,937-kg (66,000-pound) lightweight tank by another 3,401.9 kg (7,500 pounds). The 26,560-kg (58,500-pound) super lightweight tank was first flown on Space Transportation System (STS)-91 (1998), opening the door for the shuttle to deliver the heavier components needed for construction of the International Space Station.



Aerodynamics and Flight Dynamics

Introduction

Aldo Bordano

Aeroscience Challenges

Gerald LeBeau

Pieter Buning

Peter Gnoffo

Paul Romere

Reynaldo Gomez

Forrest Lumpkin

Fred Martin

Benjamin Kirk

Steve Brown

Darby Vicker

Ascent Flight Design

Aldo Bordano

Lee Bryant

Richard Ulrich

Richard Rohan

Re-entry Flight Design

Michael Tigges

Richard Rohan

Boundary Layer Transition

Charles Campbell

Thomas Horvath

The shuttle vehicle was uniquely winged so it could reenter Earth's atmosphere and fly to assigned nominal or abort landing strips. The wings allowed the spacecraft to glide and bank like an airplane during much of the return flight phase. This versatility, however, did not come without cost. The combined ascent and re-entry capabilities required a major government investment in new design, development, verification facilities, and analytical tools. The aerodynamic and flight control engineering disciplines needed new aerodynamic and aerothermodynamic physical and analytical models. The shuttle required new adaptive guidance and flight control techniques during ascent and re-entry. Engineers developed and verified complex analysis simulations that could predict flight environments and vehicle interactions. The shuttle design architectures were unprecedented and a significant challenge to government laboratories, academic centers, and the aerospace industry. These new technologies, facilities, and tools would also become a necessary foundation for all post-shuttle spacecraft developments. The following section describes a US legacy unmatched in capability and its contribution to future spaceflight endeavors.

Aeroscience Challenges

One of the first challenges in the development of the Space Shuttle was its aerodynamic design, which had to satisfy the conflicting requirements of a spacecraft-like re-entry into the Earth's atmosphere where blunt objects have certain advantages, but it needed wings that would allow it to achieve an aircraft-like runway landing. It was to be the first winged vehicle to fly through the hypersonic speed regime, providing the first real test of experimental and theoretical technology for high-speed flight. No design precedents existed to help establish necessary requirements. The decision that the first flight would carry a crew further complicated the challenge. Other than approach and landing testing conducted at Dryden Flight Research Center, California, in 1977, there would be no progressive "envelope" expansion as is typically done for winged aircraft. Nor would there be successful uncrewed launch demonstrations as had been done for all spacecraft preceding the shuttle. Ultimately, engineers responsible for characterizing the aeroscience environments for the shuttle would find out if their collective predictions were correct at the same moment as the rest of the world: during the launch and subsequent landing of Space Transportation System (STS)-1 (1981).

Aeroscience encompasses the engineering specialties of aerodynamics and aerothermodynamics. For the shuttle, each specialty was primarily associated with analysis of flight through the Earth's atmosphere.

Aerodynamics involves the study of local pressures generated over the vehicle while in flight and the resultant integrated forces and



Early conceptual designs for the Orbiter looked much like a traditional airplane with a fairly sharp nose, straight wings, and common horizontal and vertical stabilizers, as shown in this artist's rendering. As a result of subsequent aerodynamic and aerothermodynamic testing and analysis, NASA made the nose more spherical to reduce heating and used a double delta wing planform due to the severe heating encountered by straight wings and the horizontal stabilizer.

moments that, when coupled with forces such as gravity and engine thrust, determine how a spacecraft will fly. Aerothermodynamics focuses on heating to the spacecraft's surface during flight. This information is used in the design of the Thermal Protection System that shields the underlying structure from excessive temperatures. The design of the shuttle employed state-of-the-art aerodynamic and aerothermodynamic prediction techniques of the day and subsequently expanded them into previously uncharted territory.

The historical precedent of flight testing is that it is not possible to "validate"—or prove—that aerodynamic predictions are correct until vehicle performance is measured at actual flight conditions. In the case of the shuttle, preflight predictions needed to be accurate enough to establish sufficient confidence to conduct the first orbital

flight with a crew on board. This dictated that the aerodynamic test program had to be extremely thorough. Further complicating this goal was the fact that much of the expected flight regime involved breaking new ground, and thus very little experimental data were available for the early Space Shuttle studies.

Wind tunnel testing—an experimental technique used to obtain associated data—forces air past a scaled model and measures data of interest, such as local pressures, total forces, or heating rates. Accomplishing the testing necessary to cover the full shuttle flight profile required the cooperation of most of the major wind tunnels in North America. The Space Shuttle effort was the largest such program ever undertaken by the United States. It involved a traditional phased approach in the programmatic design evolution of the shuttle configuration.

The shuttle started on the launch pad composed of four primary aerodynamic elements: the Orbiter; External Tank; and two Solid Rocket Boosters (SRBs). It built speed as it rose through the atmosphere. Aeronautical and aerospace engineers often relate to speed in terms of Mach number—the ratio of the speed of an object relative to the speed of sound in the gas through which the object is flying. Anything traveling at less than Mach 1 is said to be subsonic and greater than Mach 1 is said to be supersonic. The flow regime between about Mach 0.8 and Mach 1.2 is referred to as being transonic.



This photo shows clouds enveloping portions of the vehicle (STS-34 [1987]) during ascent. When the launch vehicle was in the transonic regime, shocks formed at various positions along the vehicle to recompress the flow, which greatly impacted the structural loads and aerodynamics. Such shocks, which abruptly transition the flow from supersonic to subsonic flow, were positioned at the trailing edge of the condensation “clouds” that could be seen enveloping portions of the vehicle during ascent. These clouds were created in localized areas of the flow where the pressure and temperature conditions caused the ambient moisture to condense.


Aerodynamic loads decreased to fairly low levels as the shuttle accelerated past about Mach 5 and the atmospheric density decreased with altitude, thus the aerodynamic testing for the ascent configuration was focused on the subsonic through high supersonic regimes.

Other aspects of the shuttle design further complicated the task for engineers. Aerodynamic interference existed between the shuttle’s four elements and altered the resultant pressure loads and aerodynamics on neighboring elements. Also, since various shuttle elements were designed to separate at different points in the trajectory, engineers had to consider the various relative positions of the elements during separation. Yet another complication was the effect of plumes generated by SRBs and Space Shuttle


Main Engines (SSMEs). The plume flow fields blocked and diverted air moving around the spacecraft, thus influencing pressures on the aft surfaces and altering the vehicle’s aerodynamic characteristics.

Unfortunately, wind tunnel testing with gas plumes was significantly more expensive and time consuming than “standard” aerodynamic testing. Thus, the approach implemented was to use the best available testing techniques to completely characterize the basic “power-off” (i.e., no plumes) database. “Power-on” (i.e., with plumes) effects were then measured from a limited number of exhaust plume tests and added to the power-off measurements for the final database.

The re-entry side of the design also posed unique analysis challenges. During ascent, the spacecraft continued



While it may be intuitive to include the major geometric elements of the launch vehicle (Orbiter, External Tank, and two Solid Rocket Boosters) in aerodynamic testing, it was also important to include the plumes emanating from the three main engines on the Orbiter as well as the boosters. The tests were conducted in the 4.9-m (16-ft) Transonic Wind Tunnel at the US Air Force Arnold Engineering and Development Center, Tennessee.



Every effort was made to accurately predict a vehicle's aerodynamic characteristics using wind tunnel testing. Engineers also had to be aware of anything that could adversely affect the results. This image is of the NASA Ames Research Center 2.4 x 2.1 m (8 x 7 ft) Unitary Wind Tunnel, California.

to accelerate past the aerodynamically relevant portion of the ascent trajectory. During re-entry, this speed was carried deep into the atmosphere until there was sufficient atmospheric density to measurably dissipate the related kinetic energy. Therefore, the aerodynamics of the Orbiter were critical to the design of the vehicle from speeds as high as Mach 25 down through the supersonic and subsonic regimes to landing, with the higher Mach numbers being characterized by complex physical gas dynamics that greatly influenced the aerodynamics and heating on the vehicle compared to lower supersonic Mach numbers.

Challenges associated with wind tunnel testing limited direct applicability to the actual flight environment that engineers were interested in simulating, such as: subscale modeling of the vehicle necessary to fit in the wind tunnel and

the effect on flow-field scaling; the support structure used to hold the aerodynamic model in the wind tunnel test section, which can affect the flow on the model itself; and any influence of the wind tunnel walls. To protect against any inaccuracies in the database, each aerodynamic coefficient was additionally characterized by an associated uncertainty. Great care had to be taken to not make the uncertainties too large due to the adverse effect an uncertainty would have on the design of the flight control system and the ultimate performance of the spacecraft.

In the end, given the 20,000 hours of wind tunnel test time consumed during the early design efforts and the 80,000 hours required during the final phases, a total of 100,000 hours of wind tunnel testing was conducted for aerodynamic, aerothermodynamic, and structural dynamic testing to characterize the various shuttle system elements.

Initial Flight Experience

Traditionally, a flight test program was used to validate and make any necessary updates to the preflight aerodynamic database. While flight test programs use an incremental expansion of the flight envelope to demonstrate the capabilities of an aircraft, this was not possible with the shuttle. Once launched, without initiation of an abort, the shuttle was committed to flight through ascent, orbital operations, re-entry, and landing. NASA placed a heavy emphasis on comparison of the predicted vehicle performance to the observed flight performance during the first few shuttle missions, and those results showed good agreement over a majority of flight regimes.

Two prominent areas, however, were deficient: predictions of the launch vehicle's ascent performance, and the "trim" attitude of the Orbiter during the early phase of re-entry.

On STS-1, the trajectory was steeper than expected, resulting in an SRB separation altitude about 3 km (1.9 miles) higher than predicted. Postflight analysis revealed differences between preflight aerodynamic predictions and actual aerodynamics observed by the shuttle elements due to higher-than-predicted pressures on the shuttle's aft region. It was subsequently determined that wind tunnel predictions were somewhat inaccurate because SRB and SSME plumes were not adequately modeled. This issue also called into question the structural assessment of the wing, given the dependence on the preflight prediction of aerodynamic loads. After additional testing and cross checking with flight data, NASA was able to verify the structural assessment.



The Space Shuttle Enterprise was used to conduct approach and landing testing (1977) at the Dryden Flight Research Center, California. In the five free flights, the astronaut crew separated the spacecraft from the Shuttle Carrier Aircraft and maneuvered to a landing. These flights verified the Orbiter's pilot-guided approach and landing capability and verified the Orbiter's subsonic airworthiness in preparation for the first crewed orbital flight.

Another discrepancy occurred during the early re-entry phase of STS-1. Nominally, the Orbiter was designed to reenter in an attitude with the nose of the vehicle inclined 40 degrees to the oncoming air. In aeronautical terms, this is a 40-degree angle of attack. To aerodynamically control this attitude, the Orbiter had movable control surfaces on the trailing edge of its wings and a large “body flap.” To maintain the desired angle of attack, the Orbiter could adjust the position of the body flap up out of the flow or down into the flow, accordingly. During STS-1, the body flap deflection was twice the amount than had been predicted would be required and was uncomfortably close to the body flap’s deployment limit of 22.5 degrees. NASA determined that the cause was “real gas effects”—a phenomenon rooted in high-temperature gas dynamics.

During re-entry, the Orbiter compressed the air of the atmosphere as it smashed into the atmosphere at hypersonic speed, causing the temperature of the air to heat up thermodynamically. The temperature rise was so extreme that it broke the chemical bonds that hold air molecules together, fundamentally altering how the flow around the Orbiter compressed and expanded. These high-temperature gas dynamic effects influenced the pressure distribution on the aft portion of the heat shield, thus affecting its nominal trim condition. The extent to which this effect affected the Orbiter had not been observed before; thus, it was not replicated in the wind tunnel testing used during the design phase. NASA researchers developed an experimental technique to simulate this experience using a special test gas that mimicked the behavior of high-temperature air at the lower temperatures achieved during wind tunnel testing.

Advances in Computational Aerosciences

The use of computational fluid dynamics was eventually developed as a complementary means of obtaining aeroscience information. Engineers used computers to calculate flow-field properties around the shuttle vehicle for a given flight condition. This included pressure, shear stress, or heating on the vehicle surface, as well as density, velocity, temperature, and pressure of the air away from the vehicle. This was accomplished by numerically solving a complex set of nonlinear partial differential equations that described the motion of the fluid and satisfied a fundamental requirement for conservation of mass, momentum, and energy everywhere in the flow field.

Given its relative lack of sophistication and maturity, coupled with the modest computational power afforded by computers in the 1970s, computational fluid dynamics played almost no role in the development of the Space Shuttle aerodynamic database. In the following decades, bolstered by exponential increases in computer capabilities and continuing research, computational fluid dynamics took on a more prominent role. As with any tool, demonstrated validation of results with closely related experimental or flight data was an essential step prior to its use.


The most accurate approach for using wind tunnel data to validate computational fluid dynamics predictions was to directly model the wind tunnel as closely as possible, computationally. After results were validated at wind tunnel conditions, the computational fluid dynamics tool could be run at the flight conditions and used directly, or the difference between the computed flight and

wind tunnel predictions could be added to the baseline experimental wind tunnel measured result.

Because different flight regimes have unique modeling challenges, NASA developed separate computational fluid dynamics tools that were tuned to specific flight regimes. This allowed the computational algorithms employed to be optimized for each regime. Although not available during the preflight design of the Space Shuttle, several state-of-the-art computational tools were created that contributed significantly to the subsequent success of the shuttle, providing better understanding of control surface effectiveness, aerodynamic interference effects, and damage assessment. The examples of OVERFLOW and Langley Aerothermodynamic Upwind Relaxation Algorithm (LAURA) software packages were both based on traditional computational fluid dynamics methods while the digital to analog converter (DAC) software employed special-purpose algorithms that allowed it to simulate rarefied, low-density flows.

The OVERFLOW computational fluid dynamics tool was optimized for lower Mach number subsonic, transonic, and supersonic flows. It was thus most applicable for ascent and late re-entry simulations. Additionally, its underlying methodology was based on an innovative and extremely flexible approach for discretization of the domain around the vehicle. This was especially beneficial for analysis of a complex geometry like the shuttle.

The development of this computational fluid dynamics tool allowed engineers to effectively model the requisite geometric detail of the launch vehicle, as well as the plumes. OVERFLOW was subsequently used to investigate



This image depicts the geometric detail included in this high-fidelity modeling capability, as well as some representative results produced by the OVERFLOW tool. The OVERFLOW computational fluid dynamics tool was optimized for lower Mach number subsonic, transonic, and supersonic flows. The surface pressure is conveyed by a progressive color scale that corresponds to the pressure magnitude. A similar color scale with a different range is used to display Mach number in the flow field. OVERFLOW provided extremely accurate predictions for the launch vehicle aerodynamic environments. Color contouring depicts the nominal heating distribution on the Orbiter, where hotter colors represent higher values and cooler colors represent lower values.

the effect of design changes to the shuttle's aerodynamic performance. Some of these directly impacted shuttle operations, including all of the changes made to the tank after the Columbia accident in 2003 to help minimize the debris. Additionally, OVERFLOW solutions became a key element in the program's risk assessment for ascent debris, as the detailed flow-field


information it provided was used to predict trajectories of potential debris sources. OVERFLOW became a key tool for commercial and military transport analyses and was heavily used by industry as well as other NASA programs.

The LAURA package was another traditional computational fluid


dynamics code, but designed specifically to predict hypersonic flows associated with re-entry vehicles. It incorporated physical models that account for chemical reactions that take place in air at the extremely high temperatures produced as a spacecraft reenters an atmosphere, as well as the temporal speed at which these reactions take place. This was essential, as the “resident” time a fluid element was in the vicinity of the Orbiter was extremely short given that the vehicle traveled more than 20 times the speed of sound and the chemical reactions taking place in the surrounding fluid occurred at a finite rate.

LAURA underwent extensive validation through comparisons to a wide body of experimental and flight data, and it was also used to investigate, reproduce, and answer questions associated with the Orbiter body flap trim anomaly. LAURA was used extensively during the post-Columbia accident investigation activities and played a prominent role in supporting subsequent shuttle operations. This included assessing damaged or repaired Orbiter Thermal Protection System elements, as well as providing detailed flow field characteristics. These characteristics were assessed to protect against dangerous early transitioning of the flow along the heat shield of the Orbiter from smooth laminar flow to turbulent conditions, and thus

NASA used the Direct Simulation Monte Carlo method to simulate low-density flows, such as those created by maneuvering thrusters during orbital rendezvous and docking of the shuttle to the space station. While the method made use of a distinctly different modeling technique to make its predictions, it produced the same detailed information about the flow field as would a traditional computational fluid dynamics technique.



Special computational fluid dynamics programs appropriately model the complex chemically reacting physics necessary to accurately predict a spacecraft's aerodynamic characteristics and the aerothermodynamic heating it will experience. Heating information was needed to determine the appropriate materials and thickness of the Thermal Protection System that insulated the underlying structure of the vehicle from hot gases encountered during re-entry into Earth's atmosphere. Color contouring depicts the nominal heating distribution on the Orbiter, where hotter colors represent higher values and cooler colors represent lower values.



greatly elevated heating that would have endangered the vehicle and crew.

While traditional computational fluid dynamics tools proved extremely useful, their applicability was limited to denser portions of the atmosphere. NASA recognized the need to also be able to perform accurate analysis of low-density flows. Subsequently, the agency invested in the development of a state-of-the-art computer program that would be applicable to low-density rarefied flows. This program was based on the Direct Simulation Monte Carlo (DSMC) method—which is a simulation of a gas at the molecular level that tracks molecules through physical space and their subsequent deterministic collisions with a surface and representative collisions with other molecules. The resulting software, named the DSMC Analysis Code, was used extensively in support of shuttle missions to the Russian space station Mir and the International Space Station, as well as Hubble Space Telescope servicing missions. It also played a critical role in the analysis of the Mars Global Surveyor (1996) and the Mars Odyssey (2001) missions.

Leveraging the Space Shuttle Experience

Never before in the history of flight had such a complex vehicle and challenging flight regime been characterized. As a result of this challenge, NASA developed new and improved understanding of the associated physics, and subsequently techniques and tools to more accurately simulate them. The aerospace techniques and technologies that successfully supported the Space Shuttle are useful for exploration of our solar system.

Ascent Flight Design


NASA's challenge was to put wings on a vehicle and have that vehicle survive the atmospheric heating that occurred during re-entry into Earth's atmosphere. The addition of wings resulted in a much-enhanced vehicle with a lift-to-drag ratio that allowed many abort options and a greater cross-range capability, affording more return-to-Earth opportunities. This Orbiter capability did, however, create a unique ascent flight design challenge. The launch configuration was no longer a smooth profiled rocket. The vehicle during ascent required new and complex aerodynamic and structural load relief capabilities.

The Space Shuttle ascent flight design optimized payload to orbit while operating in a constrained environment. The Orbiter trajectory needed to restrict wing and tail structural loading during maximum dynamic pressure

and provide acceptable first stage performance. This was achieved by flying a precise angle of attack and sideslip profile and by throttling the main engines to limit dynamic pressure to five-times-gravity loads. The Solid Rocket Boosters (SRBs) had a built-in throttle design that also minimized the maximum dynamic pressure the vehicle would encounter and still achieve orbital insertion.


During the first stage of ascent, the vehicle angle of attack and dynamic pressure produced a lift force from the wings and produced vehicle structural loading. First stage guidance and control algorithms ensured that the angle of attack and sideslip did not vary significantly and resulted in flying through a desired keyhole. The keyhole was defined by the product of dynamic pressure and angle of attack. The product of dynamic pressure and sideslip maintained the desired loading on the vehicle tail.

Varying Throttle to Meet Dynamic Pressure Constraints During Ascent



During ascent, the shuttle's main engines were throttled down due to dynamic pressure constraints. The goal was to get as close as possible to the constraints to maximize performance.

Flying Through a Keyhole



Load dispersions, which are mostly due to atmospheric and thrust variations, added further constraints to the shuttle's flight. To avoid the various load dispersions at certain Mach numbers, the shuttle had to deviate from its optimum angle of attack.

Because day-of-launch winds aloft significantly altered vehicle angle of attack and sideslip during ascent, balloon measurements were taken near liftoff and in proximity of the launch site. Based on these wind measurements, Orbiter guidance parameters were biased and updated via telemetry.

Also during first stage, a roll maneuver was initiated after the vehicle cleared the tower. This roll maneuver was required to achieve the desired orbital inclination and put the vehicle in a heads-down attitude during ascent.

Vehicle performance was maximized during second stage by a linear steering law called powered explicit guidance. This steering law guided the vehicle to orbital insertion and provided abort capability to downrange abort sites or return to launch site. Ascent performance was maintained. If one main engine failed, an intact abort could be achieved to a safe landing site. Such aborts allow the Orbiter and crew to either fly at a lower-than-planned orbit or land.

Ascent flight design was also constrained to dispose the External Tank (ET) in safe waters—either the Indian Ocean or the Pacific Ocean—or in a location where tank debris was not an issue.

After main engine cutoff and ET separation, the remaining main engine fuel and oxidizer were dumped. This event provided some additional performance capability.

After the shuttle became operational, additional ascent performance was added to provide safe orbit insertion for some heavy payloads. Many guidance and targeting algorithm additions provided more payload capability. For example, standard targets were replaced by direct targets, resulting in one Orbital Maneuvering System maneuver instead of two. This saved propellant and resulted in more payload to orbit.

The ascent flight design algorithms and techniques that were generated for the shuttle will be the foundation for ascent flight of any new US launch vehicle.



Ascent Abort

During ascent, a first stage Orbiter main engine out required the shuttle to return to the launch site. The on-board guidance adjusted the pitch profile to achieve SRB staging conditions while satisfying structural and heating constraints. For a side Orbiter main engine out, the vehicle was rolled several degrees so that the normal aerodynamic force canceled the side force induced by the remaining good side engine. Also, vehicle sideslip was maintained near zero to satisfy structural constraints.

After the SRBs were safely separated, second stage guidance commanded a fixed pitch attitude around 70 degrees to minimize vehicle heating and burn the fuel no longer required. This was called the fuel dissipation phase and lasted until approximately 2% of the fuel remained. At this point, guidance commanded the vehicle to turn around and fly back to the launch site using the powered explicit guidance algorithm. As the vehicle returned, it was pitched down so the ET could be safely separated. Dynamic pressure was also minimized so a safe re-entry could occur.

During second stage ascent, a main engine failure usually required the vehicle to abort to a transatlantic landing site. An abort to a downrange landing site was preferred to a return to launch site to reduce complex trajectory targeting and minimize the loads and heating environments, therefore increasing abort success. If a main engine failure occurred late during second stage, an abort to a safe orbit was possible. Abort to orbit was preferred over an abort to a transatlantic landing site. Once the shuttle was in a safe orbit, the vehicle could perform a near nominal re-entry and return to the planned US landing strip.

Space Shuttle Ascent Abort Scenarios



The shuttle had four types of intact aborts: Return to Launch Site; Transatlantic Abort Landing; Abort to Orbit; and Abort Once Around. The aborts are presented as they occurred in the mission timeline. The preferred order of selecting aborts based on performance and safety was: Abort to Orbit; Abort Once Around; Transatlantic Abort Landing; and Return to Launch Site.

If more than one main engine failed during ascent, a contingency abort was required. If a contingency abort was called during first stage, guidance would pitch the vehicle up to loft the trajectory, thereby minimizing dynamic pressure and allowing safe separation of the SRBs and ET. After these events, a pullout maneuver would be performed to bring the vehicle to a gliding flight so a crew bailout could occur.

Two engines out early during second stage allowed the crew to attempt a landing along the US East Coast at predefined landing strips. Two engines out late in second stage allowed an abort to a transatlantic site or abort to safe orbit, depending on the time of the second failure.

In general, Mission Control used vehicle telemetry and complex vehicle performance predictor algorithms to assist the crew in choosing the best abort guidance targets and a safe landing site. The Abort Region Determinator was the primary ground flight design tool that assisted Mission Control in making abort decisions. If communication with the ground was lost, the crew would use on-board computer data and cue cards to assist in selecting the abort mode.

Summary

The shuttle ascent and ascent flight design were complex. NASA developed and verified many innovative guidance algorithms to accomplish mission objectives and maintain vehicle and crew safety. This legacy of flight techniques and computer tools will prove invaluable to all new spacecraft developments.


Re-entry Flight Design

The shuttle vehicle reentered the Earth's atmosphere at over 28,000 km per hour (kph) (17,400 mph)—about nine times faster than the muzzle speed of an M16 bullet. Designing a guidance system that safely decelerated this rapidly moving spacecraft to runway landing speeds while respecting vehicle and crew constraints was a daunting challenge, one that the shuttle re-entry guidance accomplished.


The shuttle re-entry guidance provided steering commands from

initial re-entry at a speed of 28,000 kph (17,400 mph), an altitude of 122 km (76 miles), and a distance of 7,600 km (4,722 miles) from the runway until activation of terminal area guidance (a distance of about 90 km [56 miles] and 24 km [15 miles] altitude from the runway). During this interval, a tremendous amount of kinetic energy was transferred into heat energy as the vehicle slowed down. This was all done while the crew experienced only about 1.5 times the acceleration of gravity (1.5g). As a comparison, 1g acceleration is what we feel while sitting on a chair at sea level.

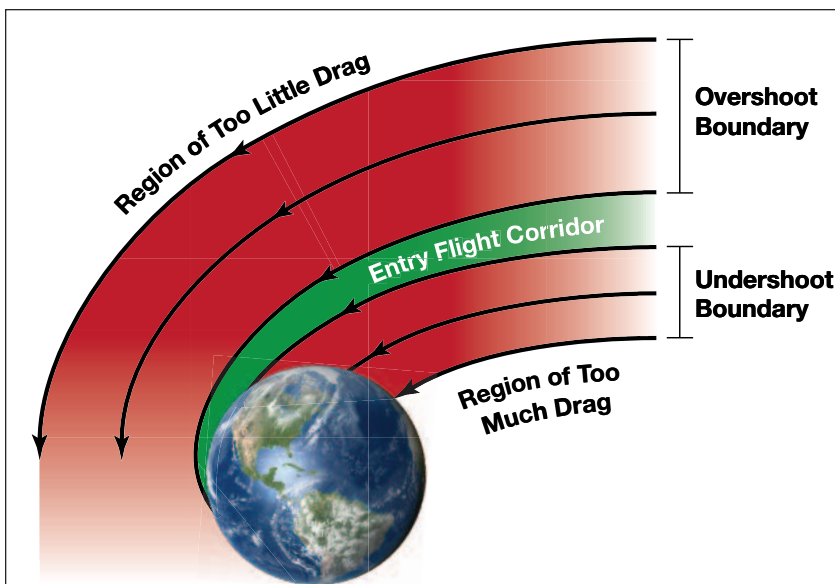
Entry Guidance Drag Velocity Profile



Shuttle re-entry guidance was segmented into several phases—each designed to satisfy unique constraints during flight. The narrow region of acceptable flight conditions was called the “flight corridor.” The surface temperature constraints resided at the lower altitude and high drag “undershoot” side of the flight corridor. In contrast, if the vehicle flew too close to the “overshoot” boundary, it would not have enough drag acceleration to reach the landing site and could possibly skip back into orbit. As the vehicle penetrated deeper into the atmosphere, the undershoot corridor was redefined by the vehicle control system and dynamic pressure constraints.



Shuttle re-entry guidance generated bank angle and angle-of-attack commands. The body flap was used to control the angle of attack by balancing the aerodynamic forces and moments about the vehicle center of gravity. The bank angle controlled the direction of the lift vector about the wind velocity vector at a fixed angle of attack. Drag, which was opposite to the wind-relative velocity, slowed the vehicle down. Lift was normal to the drag vector and was used to change the rate at which the vehicle reentered the atmosphere. The total normal load force was the sum of the lift acceleration and drag acceleration and resulted in the force felt by the crew.



The Entry Flight Corridor defined the atmospheric re-entry angles required for safe re-entry flight. Before any successful re-entry from low-Earth orbit could occur, the shuttle needed to fire engines to place the vehicle on a trajectory that intercepted the atmosphere. This deorbit maneuver had to be executed precisely. With too steep of a re-entry, the guidance could not compute steering commands that would stop the vehicle from overheating. With too shallow of a re-entry, the guidance could not adequately control the trajectory or, for very shallow trajectories, even stop the vehicle from skipping back out into space. The area between these two extremes was called the Entry Flight Corridor.

How did Space Shuttle Guidance Accomplish This Feat?

First, it's important to understand how the shuttle was controlled. Air molecules impacting the vehicle's surface imparted a pressure or force over the vehicle's surface. The shuttle used Reaction Control System jets initially to control the attitude of the vehicle; however, as the dynamic pressure increased on entering denser atmosphere, the position of the body flap was used to control the angle of attack and the ailerons were used to control bank.

Changing the angle of attack had an immediate effect on the drag acceleration of the vehicle, whereas changing the bank angle had a more gradual effect. It took time for the vehicle to decelerate into different portions of the atmosphere where density and speed affected drag. Controlling the direction of the vehicle lift vector by banking the vehicle was the primary control mechanism available to achieve the desired landing target. The vehicle banked about the relative velocity vector using a combination of aft yaw Reaction Control System jets and aileron deflection. The lift vector moved with the vehicle as it banked about the wind vector. The angle of attack was maintained constant during these maneuvers by the balanced aerodynamic forces at a given body flap trim position. The vehicle banked around this wind vector, keeping the blunt side of the shield facing against the flow of the atmosphere. Banking about the wind vector until the lift pointing down accelerated the vehicle into the atmosphere. Over time, this increased drag caused the vehicle to decelerate quickly. Banking about the wind vector until the lift vector pointed up accelerated the vehicle out of the

atmosphere. Over time, this decreased the drag acceleration and caused the vehicle to decelerate gradually. Control of the vehicle lift-and-drag acceleration by bank angle and angle-of-attack modulation were the two primary control parameters used to fly the desired range and cross range during re-entry. These concepts had to be clearly grasped before it was possible to understand the operation of the guidance algorithm.


Within each guidance phase, it was possible to use simple equations to analytically compute how much range was flown. As long as the shuttle trajectory stayed “close” to reference profiles, the guidance algorithm could analytically predict how far the vehicle would fly.

By piecing together all of the guidance segments, the total range flown from the current vehicle position all the way to the last guidance phase could be predicted and compared to the actual range required to reach the target. Any difference between the analytically computed range and the required range would trigger an adjustment in the drag-velocity/energy references to remove that range error. The analytic reference profiles were computed every guidance step (1.92 seconds) during flight. In this manner, any range error caused by variations in the environment, navigated state, aerodynamics, or mass properties was sensed and compensated for with adjustments to the real-time computed drag-velocity or drag-energy reference profiles.

In fact, the entire shuttle re-entry guidance system could be described as a set of interlocked drag-velocity or drag-energy pieces that would fly the required range to target and maintain the constraints of flight.

Boundary Layer Transition

Accurate characterization of the aerothermodynamic heating experienced by a spacecraft as it enters an atmosphere is of critical importance to the design of a Thermal Protection System. More intense heating typically requires a thicker Thermal Protection System, which increases a vehicle's weight. During the early phase of entry, the flow near the surface of the spacecraft—referred to as the boundary layer—has a smooth laminar profile. Later in the trajectory, instabilities develop in the boundary layer that cause it to transition to a turbulent condition that can increase the heating to the spacecraft by up to a factor of 4 over the laminar state. Subsequently, a Boundary Layer Transition Flight Experiment was conceived and implemented on Space Shuttle Discovery's later flights. This experiment employed a fixed-height protuberance (speed bump) on the underside of the wing to perturb and destabilize the boundary layer. NASA used instrumentation to measure both the elevated heating on the protuberance as well as the downstream effect so that the progression of the transition could be captured. The experiment provided foundational flight data that will be essential for the validation of future ground-based testing techniques or computational predictions of this flow phenomenon, thus helping improve the design of all future spacecraft.



A NASA team—via a US Navy aircraft—captured high-resolution, calibrated infrared imagery of Space Shuttle Discovery's lower surface in addition to discrete instrumentation on the wing, downstream, and on the Boundary Layer Transition Flight Experiment protuberance. In the image, the red regions represent higher surface temperatures.

Constant Heat-rate Phase

The guidance phase was required to protect the structure and interior from the blast furnace of plasma building


up outside of the vehicle. That blast furnace was due to the high-velocity impact of the vehicle with the air in the atmosphere.

The Thermal Protection System surface was designed to withstand extremely high temperatures before the temperature limits of the material were exceeded. Even after a successful landing, structural damage from heating could make the vehicle un-reuseable; therefore, it was essential that the surface remain within those limits. To accomplish this, different parts of the vehicle were covered with different types of protective material, depending on local heating.

The objective of the re-entry guidance design during this phase was to ensure that the heat-rate constraints of the Thermal Protection System were not compromised. That is why the constant heat-rate phase used quadratic drag-velocity segments. A vehicle following a drag acceleration profile that was quadratic in velocity experienced a constant rate of heating on the Thermal Protection System. Because the shuttle tile system was designed to radiate heat, the quadratic profiles in shuttle guidance were designed to provide an equilibrium heating environment where the amount of heat transferred by the tiles and to the substructure was balanced by the amount of heat radiated. This meant that there was a temperature at which the radiant heat flux away from the surface matched the rate of atmospheric heating. Once the vehicle Thermal Protection System reached this equilibrium temperature, there would no longer be a net heat flow into the vehicle.

The existence of a temperature limit on the Thermal Protection System material implied the existence of a maximum heat rate the vehicle could withstand. As long as guidance commanded the vehicle to achieve a quadratic velocity reference that was at or below the surface temperature

Typical Angle-of-Attack Profile



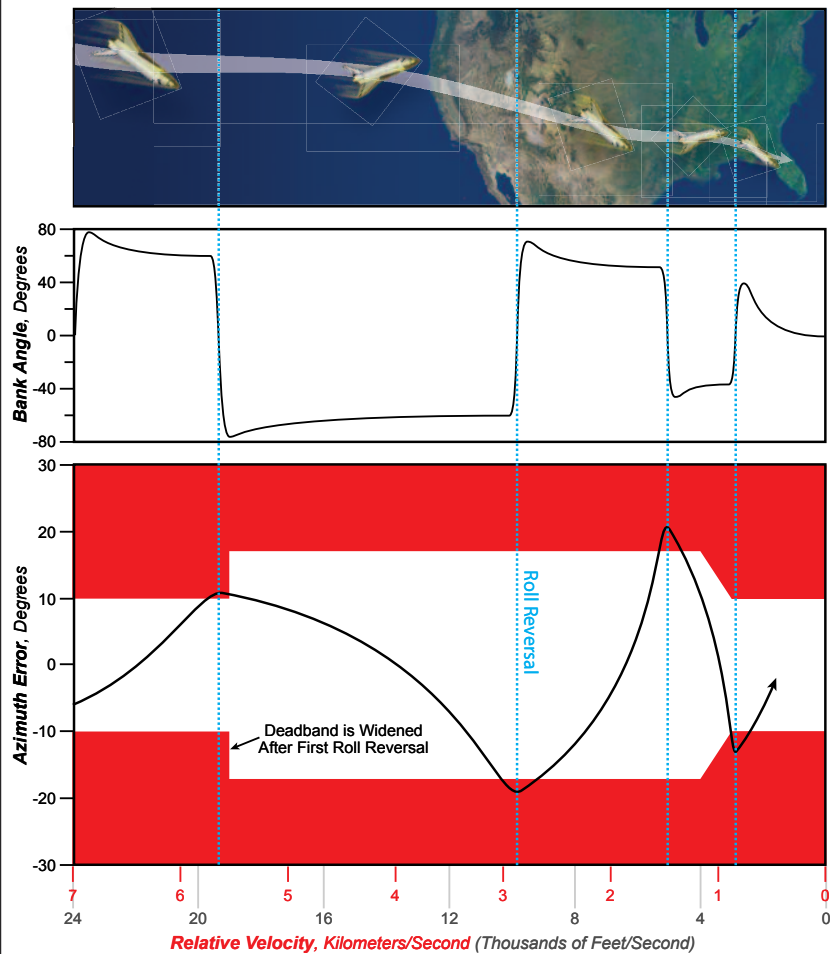
The shuttle guidance was forced to balance conflicting trades to minimize the weight, cost, and complexity of the required subsystems, maximize re-entry performance (range and cross-range capability), and maintain constraint margins. An ideal example was the selection of a constant angle-of-attack (Alpha) profile with a linear-velocity ramp transition. It was known that a high heat-rate trajectory would minimize the tile thickness required to protect the substructure. An initially high Alpha trim (40 degrees) was therefore selected to reduce Thermal Protection System mass and quickly dissipate energy. The 40-degree profile helped shape the forward center-of-gravity control boundaries and define the hypersonic static margin control limits provided by the body flap and ailerons. A linear ramp in the Alpha profile was then inserted to increase the lift-to-drag and cross-range capability and improve the static and dynamic stability of the vehicle.

constraint boundaries, the vehicle substructure was maintained at a safe temperature. The Thermal Protection System would be undamaged and reusable, and the crew would be comfortable.


During flight, if the vehicle was too close to the landing site target, the velocity and reference drag profiles were automatically shifted upward, causing an increase in the rate energy

is dissipated. The vehicle would, as a result, fly a shorter range. If the vehicle was too far away from the landing site, the combined velocity and reference drag profiles were automatically shifted downward, causing a reduction in the rate at which energy was dissipated. The vehicle would, as a result, fly a longer range.

Lateral Deadband



Azimuth Error



The Space Shuttle removed azimuth errors during flight by periodically executing roll reversals. These changes in the sign (plus or minus) of the vehicle bank command would shift the lift acceleration vector to the opposite side of the current orbit direction and slowly rotate the direction of travel back toward the desired target.

Equilibrium Glide Phase

As the speed of the shuttle dropped below about 6,200 m/s (20,500 ft/s), the constant heat-rate phase ended and the equilibrium glide phase began. This was an intermediate phase between high heating and the rapidly increasing deceleration that occurred as the vehicle penetrated deeper into the atmosphere. This phase determined the drag-velocity reference required to

balance gravitational and centrifugal forces on the vehicle. During this phase, only the reference drag profile in the equilibrium glide phase was modified to correct range errors. All future phases were left at their nominal setting. This ranging approach was designed into the shuttle re-entry guidance to reserve ranging capability. This enabled the vehicle to accommodate large navigation errors post ionization blackout (ground

communication and tracking loss due to plasma shield interference) and also change runway landing direction due to landing wind changes.

Constant Drag Phase

The constant drag phase began and the equilibrium glide phase ended when either the desired constant drag acceleration target of 10 m/s² (33 ft/s²)

occurred or the transition phase velocity of about 3,200 m/s (10,500 ft/s) was achieved.

During the constant drag phase, the drag-velocity reference was computed to maintain constant drag acceleration on the vehicle. This constrained the accelerations on the vehicle structure and crew. It also constrained maximum load accelerations for crew members confined to a sitting position during re-entry with normal accelerations directed along their spine. For the shuttle, the normal force constraint was set at 2.5g maximum; however, typical normal force operational design was set at 1.5g. The form of the drag-velocity reference during this phase was particularly simple since the drag accelerations were held constant. Operationally, shuttle guidance continued to command a high 40-degree angle of attack during this phase while the velocity was rapidly reduced and kinetic energy was rapidly removed from the vehicle. Guidance commanded higher drag levels to remove extra energy from the vehicle and to attain a target site that was closer than the nominal prediction. Guidance commanded lower drag levels to reduce the rate energy removed from the vehicle and to attain a target site that was farther away than the nominal prediction.

Transition Phase

When the velocity dropped below approximately 3,200 m/s (10,500 ft/s), the transition phase of guidance was entered and the constant drag phase was terminated. It was during this phase that the guidance system finally began to modulate the energy-vs.-drag reference to remove final

trajectory-range errors and issued a command to begin reducing the angle of attack. This pitch-down maneuver prepared the vehicle for transonic and subsonic flight. During the transition phase, the angle of attack was reduced and the vehicle transitioned from flying on the “back side” to the “front side” of the lift-to-drag (lift acceleration divided by drag acceleration) vs. angle-of-attack curve. A vehicle flying on the back side (at a higher angle of attack) was in an aerodynamic posture where increasing the angle of attack decreased the lift-to-drag. In this orientation, the drag on the vehicle was maximized and the vehicle dissipated a great deal of energy, which was highly desirable in the early phases of re-entry flight. A vehicle flying on the front side of the lift-to-drag curve (or at a lower angle of attack) was in an aerodynamic posture where increasing the angle of attack increased the lift-to-drag. In this front-side orientation, the drag was reduced and the vehicle sliced through the air more efficiently. Most airplanes fly on the front side of the lift-to-drag curve, and it was during the transition phase that shuttle guidance began commanding the vehicle to a flying orientation that mimicked the flight characteristics of an airplane.

It was also during the transition phase that the flight-path angle became significantly steeper. This happened naturally as the vehicle began to dig deeper into the atmosphere. A steeper angle was what influenced the formulation of the shuttle guidance to switch from velocity to energy as the independent variable in the reference drag formulation. The linear drag-energy reference acceleration did not use a shallow flight-path angle approximation as was done in the

previous guidance phases, and a concise closed-form solution for the range flown at higher flight-path angles was obtained. At the end of transition phase, the vehicle was about 90 km (56 miles) from the runway, flying at an altitude of 24 km (15 miles) and a speed of 750 m/s (2,460 ft/s).

Summary

At this point, the “unique” phase of re-entry required to direct the shuttle from low-Earth orbit was complete. Although other phases of guidance were initiated following the transition phase, these flight regimes were well understood and the guidance formulation was tailored directly for airplane flight.



Avionics, Navigation, and Instrumentation

Introduction

Gail Chapline

Reconfigurable Redundancy

Paul Sollock

Shuttle Single Event Upset Environment

Patrick O'Neill

Development of Space Shuttle Main Engine Instrumentation

Arthur Hill

Unprecedented Rocket Engine Fault-Sensing System

Tony Fiorucci

Calibration of Navigational Aides Using Global Positioning Computers

John Kiriazes

The Space Shuttle faced many vehicle control challenges during ascent, as did the Orbiter during on-orbit and descent operations. Such challenges required innovations such as fly-by-wire, computer redundancy for robust systems, open-loop main engine control, and navigational aides. These tools and concepts led to groundbreaking technologies that are being used today in other space programs and will be used in future space programs. Other government agencies as well as commercial and academic institutions also use these analysis tools. NASA faced a major challenge in the development of instruments for the Space Shuttle Main Engines—engines that operated at speeds, pressures, vibrations, and temperatures that were unprecedented at the time. NASA developed unique instruments and software supporting shuttle navigation and flight inspections. In addition, the general purpose computer used on the shuttle had static random access memory, which was susceptible to memory bit errors or bit flips from cosmic rays. These bit flips presented a formidable challenge as they had the potential to be disastrous to vehicle control.



Reconfigurable Redundancy— The Novel Concept Behind the World's First Two-Fault-Tolerant Integrated Avionics System

Space Shuttle Columbia successfully concluded its first mission on April 14, 1981, with the world's first two-fault-tolerant Integrated Avionics System—a system that represented a curious dichotomy of past and future technologies. On the one hand, many of the electronics components, having been selected before 1975, were already nearing technical obsolescence. On the other hand, it used what were then-emerging technologies; e.g.,

time-domain-multiplexed data buses, fly-by-wire flight control, and digital autopilots for aircraft, which provided a level of functionality and reliability at least a decade ahead of the avionics in either military or commercial aircraft. Beyond the technological “nuts and bolts” of the on-board system, two fundamental yet innovative precepts enabled and shaped the actual implementation of the avionics system. These precepts included the following:

- The entire suite of avionics functions, generally referred to as “subsystems”—data processing (hardware and software), navigation, flight control, displays and controls, communications and tracking, and electrical power distribution and control—would be programmatically and technically managed as an integrated set of subsystems. Given that new and unique types of complex hardware and software had to be developed and certified, it is difficult to overstate the role that approach played in keeping those activities on course and on schedule toward a common goal.
- A digital data processing subsystem comprised of redundant central processor units plus companion input/output units, resident software, digital data buses, and numerous remote bus terminal units would function as the core subsystem to interconnect all avionics subsystems. It also provided the means for the crew and ground to access all vehicle systems (i.e., avionics and non-avionics systems). There were exceptions to this, such as the landing gear, which was lowered by the crew via direct hardwired switches.



STS-1 launch (1981) from Kennedy Space Center, Florida. First crewed launch using two-fault-tolerant Integrated Avionics System.



Avionics System Patterned After Apollo; Features and Capabilities Unlike Any Other in the Industry

The preceding tenets were very much influenced by NASA's experience with the successful Apollo primary navigation, guidance, and control system. The Apollo-type guidance computer, with additional specialized input/output hardware, an inertial reference unit, a digital autopilot, fly-by-wire thruster control, and an alphanumeric keyboard/display unit represented a nonredundant subset of critical functions for shuttle avionics to perform. The proposed shuttle avionics represented a challenge for two principal reasons: an extensive redundancy scheme and a reliance on new technologies.

Shuttle avionics required the development of an overarching and extensive redundancy management scheme for the entire integrated avionics system, which met the shuttle requirement that the avionics system be "fail operational/fail safe"—i.e., two-fault tolerant with reaction times capable of maintaining safe computerized flight control in a vehicle traveling at more than 10 times the speed of high-performance military aircraft.

Shuttle avionics would also rely on new technologies—i.e., time-domain data buses, digital fly-by-wire flight control, digital autopilots for aircraft, and a sophisticated software operating system that had very limited application in the aerospace industry of that time, even for noncritical applications, much less for "man-rated" usage. Simply put, no textbooks were available to guide the design, development, and flight certification of those technologies

and only a modicum of off-the-shelf equipment was directly applicable.

Why Fail Operational/Fail Safe?

Previous crewed spacecraft were designed to be fail safe, meaning that after the first failure of a critical component, the crew would abort the mission by manually disabling the primary system and switching over to a backup system that had only the minimum capability to return the vehicle safely home. Since the shuttle's basic mission was to take humans and payloads safely to and from orbit, the fail-operational requirement was intended to ensure a high probability of mission success by avoiding costly, early termination of missions.

Early conceptual studies of a shuttle-type vehicle indicated that vehicle atmospheric flight control required full-time computerized stability augmentation. Studies also indicated that in some atmospheric flight regimes, the time required for a manual switchover could result in loss of vehicle. Thus, fail operational actually meant that the avionics had to be capable of "graceful degradation" such that the first failure of a critical component did not compromise the avionic system's capability to maintain vehicle stability in any flight regime.

The graceful degradation requirement (derived from the fail-operational/fail-safe requirement) immediately provided an answer to how many redundant computers would be necessary. Since the computers were the only certain way to ensure timely graceful degradation—i.e., automatic detection and isolation of an errant computer—some type of computerized majority-vote technique involving a minimum of three computers would be required to retain operational

status and continue the mission after one computer failure. Thus, four computers were required to meet the fail-operational/fail-safe requirement. That level of redundancy applied only to the computers. Triple redundancy was deemed sufficient for other components to satisfy the fail-operational/fail-safe requirement.

Central Processor Units Were Available Off the Shelf—Remaining Hardware and Software Would Need to be Developed

The next steps included: selecting computer hardware that was for military use yet commercially available; choosing the actual configuration, or architecture, of the computer(s), data bus network, and bus terminal units; and then developing the unique hardware and software to implement the world's first two-fault-tolerant avionics.

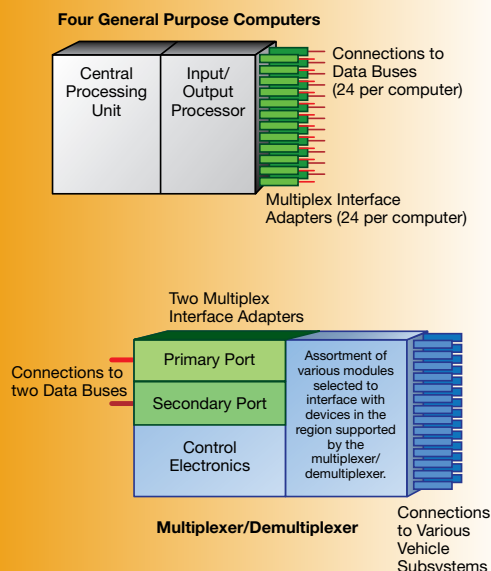
In 1973, only two off-the-shelf computers available for military aircraft offered the computational capability for the shuttle. Both computers were basic processor units—termed "central processor units"—with only minimal input/output functionality. NASA selected a vendor to provide the central processor units plus new companion input/output processors that would be developed to specifications provided by architecture designers. At the time, no proven best practices existed for interconnecting multiple computers, data buses, and bus terminal units beyond the basic active/standby manual switchover schemes.

The architectural concept figured heavily in the design requirements for the input/output processor and two other new types of hardware "boxes" as



Interconnections Were Key to Avionics Systems Success

Shuttle Systems Elements



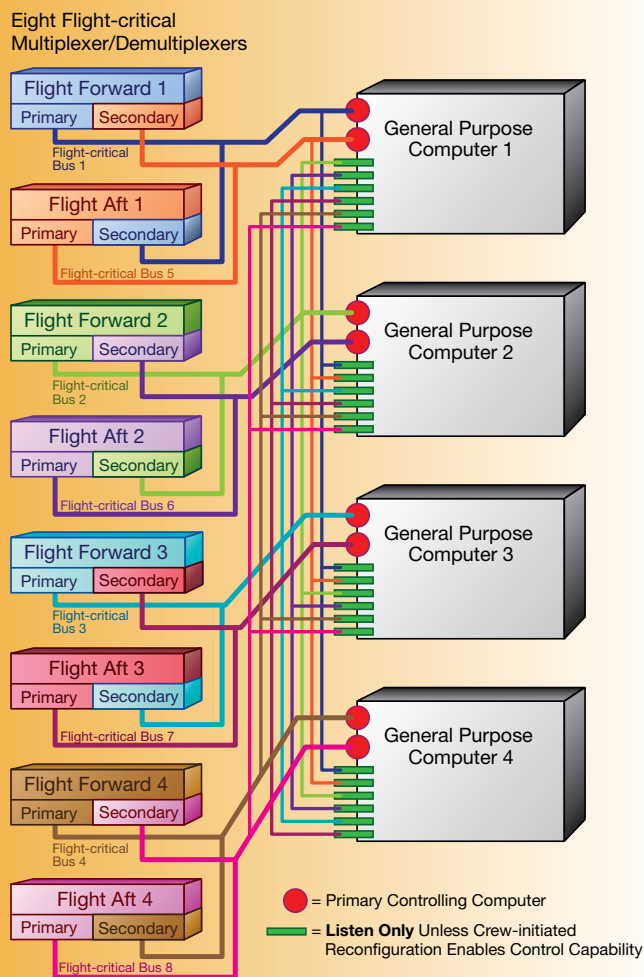
Architecture designers for the shuttle avionics system had three goals: provide interconnections between the four computers to support a synchronization scheme; provide each computer access to every data bus; and ensure that the multiplexer/demultiplexers were sufficiently robust to preclude a single internal failure from preventing computer access to the systems connected to that multiplexer/demultiplexer.

To meet those goals, engineers designed the input/output processor to interface with all 24 data buses necessary to cover the shuttle. Likewise, each multiplexer/demultiplexer would have internal

redundancy in the form of two independent ports for connections to two data buses. The digital data processing subsystem possessed eight flight-critical data buses and the eight flight-critical multiplexer/demultiplexers. They were essential to the

Shuttle Systems Redundancy

Diagram illustrates the eight "flight-critical" buses of the 24 buses on the Orbiter.



reconfiguration capability. The total complement of such hardware on the vehicle consisted of 24 data buses, 19 multiplexer/demultiplexers, and an almost equal number of other types of specialized bus terminal units.



well as the operating system software, all four of which had to be uniquely developed for the shuttle digital data processing subsystem. Each of those four development activities would eventually result in products that established new limits for the so-called “state of the art” in both hardware and software for aerospace applications.

In addition to the input/output processor, the other two new devices were the data bus transmitter/receiver units—referred to as the multiplex interface adapter—and the bus terminal units, which was termed the “multiplexer/demultiplexer.” NASA designated the software as the Flight Computer Operating System. The input/output processors (one paired with each central processor unit) was necessary to interface the units to the data bus network. The numerous multiplexer/demultiplexers would serve as the remote terminal units along the data buses to effectively interface all the various vehicle subsystems to the data bus network. Each central processor unit/input/output processor pair was called a general purpose computer.

The multiplexer/demultiplexer was an extraordinarily complex device that provided electronic interfaces for the myriad types of sensors and effectors associated with every system on the vehicle. The multiplex interface adaptors were placed internal to the input/output processors and the multiplexer/demultiplexers to provide actual electrical connectivity to the data buses. Multiplex interface adaptors were supplied to each manufacturer of all other specialized devices that interfaced with the serial data buses. The protocol for communication on those buses was also uniquely defined.

The central processor units later became a unique design for two reasons: within the first several months

in the field, their reliability was so poor that they could not be certified for the shuttle “man-rated” application; and following the Approach and Landing Tests (1977), NASA found that the software for orbital missions exceeded the original memory capacity. The central processor units were all upgraded with a newer memory design that doubled the amount of memory. That memory flew on Space Transportation System (STS)-1 in 1981.

Although the computers were the only devices that had to be quad redundant, NASA gave some early thought to simply creating four identical strings with very limited interconnections. The space agency quickly realized, however, that the weight and volume associated with so much additional hardware would be unacceptable. Each computer needed the capability to access every data bus so the system could reconfigure and regain capability after certain failures. NASA accomplished such reconfiguration by software reassignment of data buses to different general purpose computers.

The ability to reconfigure the system and regain lost capability was a novel approach to redundancy management. Examination of a typical mission profile illustrates why NASA placed a premium on providing reconfiguration capability. Ascent and re-entry into Earth’s atmosphere represented the mission phases that required automatic failure detection and isolation capabilities, while the majority of on-orbit operations did not require full redundancy when there was time to thoroughly assess the implications of any failures that occurred prior to re-entry. When a computer and a critical sensor on another string failed, the failed computer string could be reassigned via software control to a healthy computer, thereby providing a fully functional operational configuration for re-entry.

The Costs and Risks of Reconfigurable Redundancy

The benefits of interconnection flexibility came with costs, the most obvious being increased verification testing needed to certify each configuration performed as designed. Those activities resulted in a set of formally certified system reconfigurations that could be invoked at specified times during a mission. Other less-obvious costs stemmed from the need to eliminate single-point failures. Interconnections offered the potential for failures that began in one redundant element and propagated throughout the entire redundant system—termed a “single-point failure”—with catastrophic consequences. Knowing such, system designers placed considerable emphasis on identification and elimination of failure modes with the potential to become single-point failures. Before describing how NASA dealt with potential catastrophic failures, it is necessary to first describe how the redundant digital data processing subsystem was designed to function.

Establishing Synchronicity

The fundamental premise for the redundant digital data processing subsystem operation was that all four general purpose computers were executing identical software in a time-synchronized fashion such that all received the exact same data, executed the same computations, got the same results, and then sent the exact same time-synchronized commands and/or data to other subsystems.

Maintenance of synchronicity between general purpose computers was one of the truly unique features of the newly developed Flight Computer Operating System. All four general purpose computers ran in a synchronized fashion that was keyed



Shuttle Single Event Upset Environment

Five general purpose computers—the heart of the Orbiter's guidance, navigation, and flight control system—were upgraded in 1991. The iron core memory was replaced with modern static random access memory transistors, providing more memory and better performance. However, the static random access memory computer chips were susceptible to single event upsets: memory bit flips caused by high-energy nuclear particles. These single event upsets could be catastrophic to the Orbiter because general purpose computers were critical to flights since one bit flip could disable the computer.

An error detection and correction code was implemented to “fix” flipped bits in a computer word by correcting any single erroneous bit. Whenever the system experienced a memory bit flip fix, the information was downlinked to flight controllers on the ground in Houston, Texas. The event time and the Orbiter's ground track resulted in the pattern of bit flips around the Earth.

The bit flips correlated with the known space radiation environment. This phenomena had significant consequences for error detection and correction codes, which could only correct one error in a word and would be foiled by a multi-bit error. In response, system architects selected bits for each word from different chips, making it almost impossible for a single particle to upset more than one bit per word.

In all, the upgraded Orbiter general purpose computers performed flawlessly in spite of their susceptibility to ionizing radiation.



Single event upsets are indicated by yellow squares. Multi-bit single event upsets are indicated by red triangles. In these single events, anywhere from two to eight bits were typically upset by a single charged particle.

to the timing of the intervals when general purpose computers were to query the bus terminal units for data, then process that data to select the best data from redundant sensors, create commands, displays, etc., and finally output those command and status data to designated bus terminal units.

That sequence (input/process/output) repeated 25 times per second. The aerodynamic characteristics of the shuttle dictated the 25-hertz (Hz) rate. In other words, the digital autopilot had to generate stability augmentation commands at that frequency for the vehicle to retain stable flight control.

The four general purpose computers exchanged synchronization status approximately 350 times per second. The typical failure resulted in the computer halting anything resembling normal operation.



A fish-eye view of the multifunction electronic display subsystem—or “glass cockpit”—in the fixed-base Space Shuttle mission simulator at Johnson Space Center, Texas.

Early Detection of Failure

NASA designed the four general purpose computer redundant set to gracefully degrade from either four to three or from three to two members. Engineers tailored specific redundancy management algorithms for dealing with failures in other redundant subsystems based on knowledge of each subsystem’s predominant failure modes and the overall effect on vehicle performance.

NASA paid considerable attention to means of detecting subtle latent failure modes that might create the potential for a simultaneous scenario. Engineers scrutinized sensors such as gyros and accelerometers in particular for null failures. During orbital operation, the vehicle typically spent the majority of

time in a quiescent flight control profile such that those sensors were operating very near their null points. Prior to re-entry, the vehicle executed some designed maneuvers to purposefully exercise those devices in a manner to ensure the absence of permanent null failures. The respective design teams for the various subsystems were always challenged to strike a balance between early detection of failures vs. nuisance false alarms, which could cause the unnecessary loss of good devices.

Decreasing Probability of Pseudo-simultaneous Failures

There was one caveat regarding the capability to be two-fault tolerant—the system was incapable of coping with simultaneous failures since such failures obviously defeat the

majority-voting scheme. A nuance associated with the practical meaning of “simultaneous” warranted significant attention from the designers. It was quite possible for internal circuitry in complex electronics units to fail in a manner that wasn’t immediately apparent because the circuitry wasn’t used in all operations. This failure could remain dormant for seconds, minutes, or even longer before normal activities created conditions requiring use of the failed devices; however, should another unrelated failure occur that created the need for use of the previously failed circuitry, the practical effect was equivalent to two simultaneous failures.

To decrease the probability of such pseudo-simultaneous failures, the general purpose computers and multiplexer/demultiplexers were designed to constantly execute cyclic background self-test operations and cease operations if internal problems were detected.

Ferretting Out Potential Single-point Failures

Engineering teams conducted design audits using a technique known as failure modes effects analysis to identify types of failures with the potential to propagate beyond the bounds of the fault-containment region in which they originated. These studies led to the conclusion that the digital data processing subsystem was susceptible to two types of hardware failures with the potential to create a catastrophic condition, termed a “nonuniversal input/output error.” As the name implies, under such conditions a majority of general purpose computers may not have received the same data and the redundant set may have



diverged into a two-on-two configuration or simply collapsed into four disparate members.

Engineers designed and tested the topology, components, and data encoding of the data bus network to ensure that robust signal levels and data integrity existed throughout the network. Extensive laboratory testing confirmed, however, that the two types of failures would likely create conditions resulting in eventual loss of all four computers.

The first type of failure and the easiest to mitigate was some type of physical failure causing either an open or a short circuit in a data bus. Such a condition would create an impedance mismatch along the bus and produce classic transmission line effects; e.g., signal reflections and standing waves with the end result being unpredictable signal levels at the receivers of any given general purpose computer. The probability of such a failure was deemed to be extremely remote given the robust mechanical and electrical design as well as detailed testing of the hardware, before and after installation on the Orbiter.

The second type of problem was not so easily discounted. That problem could occur if one of the bus terminal units failed, thus generating unrequested output transmissions. Such transmissions, while originating from only one node in the network, would nevertheless propagate to each general purpose computer and disrupt the normal data bus signal levels and timing as seen by each general purpose computer. It should be mentioned that no amount of analysis or testing could eliminate the possibility of a latent, generic software error that could conceivably cause all

Loss of Two General Purpose Computers Tested Resilience



Space Shuttle Columbia (STS-9) makes a successful landing at Dryden Flight Research Center on Edwards Air Force Base runway, California, after reaching a fail-safe condition while on orbit.

Shuttle avionics never encountered any type (hardware or software) of single-point failure in nearly 3 decades of operation, and on only one occasion did it reach the fail-safe condition. That situation occurred on STS-9 (1983) and demonstrated the resiliency afforded by reconfiguration.

While on-orbit, two general purpose computers failed within several minutes of each other in what was later determined to be a highly improbable, coincidental occurrence of a latent generic hardware fault. By definition, the avionics was in a fail-safe condition and preparations were begun in preparation for re-entry into Earth's atmosphere. Upon cycling power, one of the general purpose computers remained failed while the other resumed normal operation. Still, with that machine being suspect, NASA made the decision to continue preparation for the earliest possible return. As part of the preparation, sensors such as the critical inertial measurement unit, which were originally assigned to the failed computer, were reassigned to a healthy one. Thus, re-entry occurred with a three-computer configuration and a full set of inertial measurement units, which represented a much more robust and safe configuration.

The loss of two general purpose computers over such a short period was later attributed to spacelight effects on microscopic debris inside certain electronic components. Since all general purpose computers in the inventory contained such components, NASA delayed subsequent flights until sufficient numbers of those computers could be purged of the suspect components.



four computers to fail. Thus, the program deemed that a backup computer, with software designed and developed by an independent organization, was warranted as a safeguard against that possibility.

This backup computer was an identical general purpose computer designed to “listen” to the flight data being collected by the primary system and make independent calculations that were available for crew monitoring. Only the on-board crew had the switches, which transferred control of all data buses to that computer, thereby preventing any “rogue” primary computers from “interfering” with the backup computer.

Its presence notwithstanding, the backup computer was never considered a factor in the fail-operational/fail-safe analyses of the primary avionics system, and—at the time of this publication—had never been used in that capacity during a mission.

Summary

The shuttle avionics system, which was conceived during the dawn of the digital revolution, consistently provided an exceptional level of dependability and flexibility without any modifications to either the basic architecture or the original innovative design concepts. While engineers replaced specific electronic boxes due to electronic component obsolescence or to provide improved functionality, they took great care to ensure that such replacements did not compromise the proven reliability and resiliency provided by the original design.

Development of Space Shuttle Main Engine Instrumentation

The Space Shuttle Main Engine operated at speeds and temperatures unprecedented in the history of spaceflight. How would NASA measure the engine’s performance?

NASA faced a major challenge in the development of instrumentation for the main engine, which required a new generation capable of measuring—and surviving—its extreme operating pressures and temperatures. NASA not only met this challenge, the space agency led the development of such instrumentation while overcoming numerous technical hurdles.

Initial Obstacles

The original main engine instrumentation concept called for compact flange-mounted transducers with internal redundancy, high stability, and a long, maintenance-free life. Challenges presented themselves immediately, however. Few instrumentation suppliers were interested in the limited market projected for the shuttle. Moreover, early engine testing disclosed that standard designs were generally incapable of surviving the harsh environments. Although the “hot side” temperatures were within the realm of jet engines, no sort of instrumentation existed that could handle both high temperatures and cryogenic environments down to minus -253°C (-423°F). Vibration environments with high-frequency spectrums extending beyond commercially testable ranges of 2,000 hertz (Hz) experienced several

hundred times the force of gravity over almost 8 hours of an engine’s total planned operational exposure. For these reasons, the endurance requirements of the instrumentation constituent materials were unprecedented.

Engine considerations such as weight, concern for leakage that might be caused by mounting bosses, and overall system fault tolerance prompted the need for greater redundancy for each transducer. Existing supplier designs, where available, were single-output devices that provided no redundancy. A possible solution was to package two or more sensors within a single transducer. But this approach required special adaptation to achieve the desired small footprint and weight.

NASA considered the option of strategically placing instrumentation devices and closely coupling them to the desired stimuli source. This approach prompted an appreciation of the inherent simplicity and reliability afforded by low-level output devices. The avoidance of active electronics tended to minimize electrical, electronic, and electromechanical part vulnerability to hostile environments. Direct mounting of transducers also minimized the amount of intermediate hardware capable of producing a catastrophic system failure response. Direct mounting, however, came at a price. In some situations, it was not possible to design transducers capable of surviving the severe environments, making it necessary to off-mount the device. Pressure measurements associated with the combustion process suffered from icing or blockage issues when hardware temperatures dropped below freezing. Purging schemes to provide positive flow in pressure tubing were necessary to alleviate this condition.



Several original system mandates were later shown to be ill advised, such as an early attempt to achieve some measure of standardization through the use of bayonet-type electrical connectors. Early engine-level and laboratory testing revealed the need for threaded connectors since the instrumentation components could not be adequately shock-isolated to prevent failures induced by excessive relative connector motion. Similarly, electromagnetic interference assessments and observed deficiencies resulted in a reconsideration of the need for cable overbraiding to minimize measurement disruption.

Problems also extended to the sensing elements themselves. The lessons of material incompatibilities or deficiencies were evident in the area of resistance temperature devices and thermocouples. The need for the stability of temperature measurements led to platinum-element

resistance temperature devices being baselined for all thermal measurements.

Aggressive engine performance and weight considerations also compromised the optimal sensor mountings. For example, it was not practical to include the prescribed straight section of tubing upstream from measuring devices, particularly for flow. This resulted in the improper loading of measuring devices, primarily within the propellant oxygen ducting. The catastrophic failure risks finally prompted the removal or relocation of all intrusive measuring devices downstream of the high-pressure oxygen turbopump. Finally, the deficiencies of vibration redline systems were overcome as processing hardware and algorithms matured to the point where a real-time synchronous vibration redline system could be adopted, providing a significant increase in engine reliability.



Weakness Detection and Solutions

In some instances, the engine environment revealed weaknesses not normally experienced in industrial or aerospace applications. Some hardware successfully passed component-level testing only to experience problems at subsystem or engine-level testing. Applied vibration spectrums mimicked test equipment limitations where frequency ranges typically did not extend beyond 2,000 Hz. The actual engine recognized no limits and continued to expose the hardware to energy above even 20,000 Hz. Therefore, a critical sensor resonance condition might only be excited during engine-level testing. Similarly, segmenting of component testing into separate vibration, thermal, and fluid testing deprived the instrumentation of experiencing the more-severe effect of combined exposures.

The shuttle's reusability revealed failure modes not normally encountered, such as those ascribed to the differences between flight and ground test environments. It was subsequently found that the microgravity exposure of each flight allowed conductive particles within instruments to migrate in a manner not experienced with units confined to terrestrial applications. Main engine pressure transducers experienced electrical shorts only during actual engine performance. During the countdown of Space Transportation System (STS)-53 (1992), a high-pressure oxidizer turbopump secondary seal measurement output pressure transducer data spike almost triggered an on-pad abort. Engineers used pressure transducers screened

Wire Failures Prompted System Redesign

High temperature measurements continued to suffer brittle fine-element wire failures until the condition was linked to operation above the material recrystallization temperature of 525°C (977°F) where excessive grain growth would result. The STS-51F (1985) in-flight engine shutdown caused by the failure of multiple resistance temperature devices mandated a redesign to a thermocouple-based system that eliminated the wire embrittlement problem.



High temperatures in some engine operating environments caused fine wires used in temperature devices to become brittle, thereby leading to failures.

© Pratt & Whitney Rocketdyne. All rights reserved.



by particle impact noise detection and microfocus x-ray examination on an interim basis until a hardware redesign could be qualified.

Effects of Cryogenic Exposure on Instrumentation

Cryogenic environments revealed a host of related material deficiencies. Encapsulating materials—necessary to provide structural support for fine wires within speed sensors—lacked resiliency at extreme low temperatures. The adverse effects of inadvertent exposure to liquefied gases within the shuttle's aft compartment produced functional failures due to excessively cold conditions. In April 1991, STS-37 was scrubbed when the high-pressure oxidizer turbopump secondary seal pressure measurement became erratic due to the damaging effects of cryogenic exposure of a circuit board.

Problems with cryogenics also extended to the externals of the instrumentation. Cryopumping—the condensation-driven pumping mechanism of inert gases such as nitrogen—severely compromised the ability of electrical connectors to maintain continuity. The normally inert conditions maintained within the engine system masked a problem with residual contamination of glassed resistive temperature devices used for cryogenic propellant measurements. Corrosive flux left over from the manufacturing process remained dormant for years until activated during extended exposures to the humid conditions at the launch site. STS-50 (1992) narrowly avoided a launch delay when a resistive temperature device had to be replaced just days before the scheduled launch date.

Expectations Exceeded

As the original main engine design life of 10 years was surpassed, part obsolescence and aging became a concern. Later designs used more current parts such as industry-standard electrical connectors. Some suppliers chose to invest in technology driven by the shuttle, which helped to ease the program's need for long-term part availability.

The continuing main engine ground test program offered the ability to use ongoing hot-fire testing to ensure that all flight hardware was sufficiently enveloped by older ground test units. Tracking algorithms and extensive databases permitted such comparisons.

Industry standards called for periodic recalibration of measuring devices. NASA excluded this from the Space Shuttle Main Engine Program at its inception to reduce maintenance for hardware not projected for use beyond 10 years. In practice, the hardware life was extended to the point that some engine components approached 40 years of use before the final shuttle flight. Aging studies validated the stable nature of instruments never intended to fly so long without recalibration.

Summary

While initial engine testing disclosed that instrumentation was a weak link, NASA implemented innovative and successful solutions that resulted in a suite of proven instruments capable of direct application on future rocket engines.

Unprecedented Rocket Engine Fault-Sensing System

The Space Shuttle Main Engine (SSME) was a complex system that used liquid hydrogen and liquid oxygen as its fuel and oxidizer, respectively. The engine operated at extreme levels of temperature, pressure, and turbine speed. At these levels, slight material defects could lead to high vibration in the turbomachinery. Because of the potential consequences of such conditions, NASA developed vibration monitoring as a means of monitoring engine health.

The main engine used both low- and high-pressure turbopumps for fuel and oxidizer propellants. Low-pressure turbopumps served as propellant boost pumps for the high-pressure turbopumps, which in turn delivered fuel and oxidizer at high pressures to the engine main combustion chamber.

The high-pressure pumps rotated at speeds reaching 36,000 rpm on the fuel side and 24,000 rpm on the oxidizer side. At these speeds, minor faults were exacerbated and could rapidly propagate to catastrophic engine failure.

During the main engine's 30-year ground test program, more than 40 major engine test failures occurred. High-pressure turbopumps were the source of a large percentage of these failures. Posttest analysis revealed that the vibration spectral data contained potential failure indicators in the form of discrete rotordynamic spectral signatures. These signatures were prime indicators of turbomachinery health and could potentially be used to mitigate




catastrophic engine failures if assessed at high speeds and in real time.

NASA recognized the need for a high-speed digital engine health management system. In 1996, engineers at Marshall Space Flight Center (MSFC) developed the Real Time Vibration Monitoring System and integrated the system into the main engine ground test program. The system used data from engine-mounted accelerometers to monitor pertinent spectral signatures. Spectral data were produced and assessed every 50 milliseconds to determine whether specific vibration amplitude thresholds were being violated.

NASA also needed to develop software capable of discerning a failed sensor from an actual hardware failure. MSFC engineers developed the sensor validation algorithm—a software algorithm that used a series of rules and threshold gates based on actual vibration spectral signature content to evaluate the quality of sensor data every 50 milliseconds.

Outfitted with the sensor validation algorithm and additional software, the Real Time Vibration Monitoring System could detect and diagnose pertinent indicators of imminent main engine turbomachinery failure and initiate a shutdown command within 100 milliseconds.

The Real Time Vibration Monitoring System operated successfully on more than 550 main engine ground tests with no false assessments and a 100% success rate on determining and disqualifying failed sensors from its vibration redlines. This, the first high-speed vibration redline system developed for a liquid engine rocket



NASA's Advanced Health Monitoring System software was integrated with the Space Shuttle Main Engine controller (shown by itself and mounted on the engine) in 2007.


© Pratt & Whitney Rocketdyne. All rights reserved.

system, supported the main engine ground test program throughout the shuttle era.

To prove that a vibration-based, high-speed engine health management system could be used for flight operations, NASA included a subscale version of the Real Time Vibration Monitoring System on Technology Flight Experiment 2, which flew on STS-96 (1999).

This led to the concept of the SSME Advanced Health Management System as a means of extending this protection to the main engine during ascent.

The robust software algorithms and redline logic developed and tested for the Real Time Vibration Monitoring System were directly applied to the Advanced Health Management System and incorporated into a redesigned



version of the engine controller. The Advanced Health Management System's embedded algorithms continuously monitored the high-pressure turbopump vibrations generated by rotation of the pump shafts and assessed rotordynamic performance every 50 milliseconds. The system was programmed to initiate a shutdown command in fewer than 120 milliseconds if vibration patterns indicated an instability that could lead to catastrophic failure.

The system also used the sensor-validation algorithm to monitor sensor quality and could disqualify a failed sensor from its redline suite or deactivate the redline altogether. Throughout the shuttle era, no other liquid engine rocket system in the world employed a vibration-based health management system that used discrete spectral components to verify safe operation.

Summary

The Advanced Health Management System, developed and certified by Pratt & Whitney Rocketdyne (Canoga Park, California) under contract to NASA, flew on numerous shuttle missions and continued to be active on all engines throughout the remainder of the shuttle flights.

Calibration of Navigational Aides Using Global Positioning Computers

The crew members awakened at 5:00 a.m. After 10 days in orbit, they were ready to return to Earth. By 7:45 a.m., the payload bay doors were closed and they were struggling into their flight suits to prepare for descent. The commander called for a weather report and advice on runway selection. The shuttle could be directed to any one of three landing strips depending on weather at the primary landing site. Regardless of the runway chosen, the descent was controlled by systems capable of automatically landing the Orbiter. The Orbiter commander took cues from these landing systems, controlled the descent, and dropped the landing gear to safely land the Orbiter. During their approach to the landing site, the Orbiter crew depended on a complex array of technologies, including a Tactical Air Navigation System and the Microwave Scanning Beam Landing System, to provide precision navigation. These systems were located at each designated landing site and had to be precisely calibrated to ensure a safe and smooth landing.

Touchdown Sites

Shuttle runways were strategically located around the globe to serve several purposes. After a routine mission, the landing sites included Kennedy Space Center (KSC) in Florida, Dryden Flight Research Center in California, and White Sands Test Facility in New Mexico. The

transoceanic abort landing sites—intended for emergencies when the shuttle lost a main engine during ascent and could not return to KSC—were located in Zaragoza and Moron in Spain and in Istres in France. Former transoceanic abort landing sites included: Dakar, Senegal; Ben Guerir, Morocco; Banjul, The Gambia; Honolulu, Hawaii; and Anderson Air Force Base, Guam. NASA certified each site.

Error Sources

Because the ground portion of the Microwave Scanning Beam Landing and Tactical Air Navigation Systems contained moving mechanical components and depended on microwave propagation, inaccuracies could develop over time that might prove detrimental to a shuttle landing. For example, antennas could drift out of mechanical adjustment. Ground settling and external environmental factors could also affect the system's accuracy. Multipath and refraction errors could result from reflections off nearby structures, terrain changes, and day-to-day atmospheric variations.

Flight inspection data gathered by the NASA calibration team could be used to determine the source of these errors. Flight inspection involved flying an aircraft through the landing system coverage area and receiving time-tagged data from the systems under test. Those data were compared to an accurate aircraft positioning reference to determine error. Restoring integrity was easily achieved through system adjustment.



Global Positioning Satellite Position Reference for Flight Inspection

Technologies were upgraded several times since first using the Global Positioning Satellite (GPS)-enabled flight inspection system. The flight inspection system used an aircraft GPS receiver as a position reference. Differences between the system under test and the position reference were recorded, processed, and displayed in real time on board the aircraft. An aircraft position reference used for flight inspection had to be several times more accurate than the system under test. Stand-alone commercial GPS systems did not have enough accuracy for this purpose. Several techniques could be used to improve GPS positioning. Differential GPS used a ground GPS receiver installed over a known surveyed benchmark. Common mode error corrections to the GPS position were calculated and broadcast over a radio data link to the aircraft. After the received corrections were applied, the on-board GPS position accuracy was within 3 m (10 ft). A real-time accuracy within 10 cm (4 in.) was achieved by using a carrier-phase technique and tracking cycles of the L-band GPS carrier signal.

NASA built several versions of the flight inspection system customized to different aircraft platforms. Different NASA aircraft were used based on aircraft availability. These aircraft include NASA's T-39 jet (Learjet), a NASA P-3 turboprop, several C-130 aircraft, and even NASA's KC-135. Each aircraft was modified with shuttle landing system receivers and antennas. Several pallets of equipment were configured and tested to reduce the installation time on aircraft to one shift.

Summary

NASA developed unique instrumentation and software supporting the shuttle navigation aids flight inspection mission. The agency developed aircraft pallets to operate, control, process, display, and archive data from several avionics receivers. They acquired and synchronized measurements from shuttle-unique avionics and aircraft platform avionics with precision time-tagged GPS position. NASA developed data processing platforms and software algorithms to graphically display and trend landing system performance in real time. In addition, a graphical pilot's display provided the aircraft pilot with runway situational awareness and visual direction cues. The pilot's display software, integrated with the GPS reference system, resulted in a significant reduction in mission flight time.

Synergy With the Federal Aviation Administration

In early 2000, NASA and the Federal Aviation Administration (FAA) entered into a partnership for flight inspection. The FAA had existing aircraft assets to perform its mission to flight-inspect US civilian and military navigation aids. The FAA integrated NASA's carrier-phase GPS reference along with shuttle-unique avionics and software algorithms into its existing control and display computers on several flight-inspection aircraft.

The NASA/FAA partnership produced increased efficiency, increased capability, and reduced cost to the government for flight inspection of the shuttle landing aids.



Software

Introduction

Gail Chapline

Steven Sullivan

Primary Software

Aldo Bordano

Geminesse Dorsey

James Loveall

Personal Computer Ground Operations Aerospace Language Offered Engineers a “View”

Avis Upton

The Ground Launch Sequencer Orchestrated Launch Success

Al Folensbee

Integrated Extravehicular Activity/Robotics

Virtual Reality Simulation

David Homan

Bradley Bell

Jeffrey Hoblit

Evelyn Miralles

Integrated Solutions for Space Shuttle Management...and Future Endeavors

Samantha Manning

Charles Hallett

Dena Richmond

Joseph Schuh

Three-Dimensional Graphics Provide Extraordinary Vantage Points

David Homan

Bradley Bell

Jeffrey Hoblit

Evelyn Miralles

Software was an integral part in the Space Shuttle hardware systems and it played a vital role in the design and operations of the shuttle. The longevity of the program demanded the on-orbit performance of the vehicle to be flexible under new and challenging environments. Because of the flexibility required, quick-turnaround training, simulations, and virtual reality tools were invaluable to the crew for new operational concepts. In addition, ground operations also benefited from software innovations that improved vehicle processing and flight-readiness testing. The innovations in software occurred throughout the life of the program. The topics in this chapter include specific areas where engineering innovations in software enabled solutions to problems and improved overall vehicle and process performance, and have carried over to the next generation of space programs.



Primary Software

NASA faced notable challenges in the development of computer software for the Space Shuttle in the early 1970s. Only two avionics computers were regarded as having the potential to perform the complex tasks that would be required of them. Even though two options existed, these candidates would require substantial modification. To further compound the problem, the 1970s also suffered a noticeable absence of off-the-shelf microcomputers. Large-scale, integrated-circuit technology had not yet reached the level of sophistication necessary for Orbiter

use. This prompted NASA to continue its search for a viable solution.

NASA soon concluded that core memory was the only reasonable choice for Orbiter computers, with the caveat that memory size was subject to power and weight limitations as well as heat constraints. The space agency still faced additional obstacles: data bus technology for real-time avionics systems was not yet fully operational; the use of tape units for software program mass storage in a dynamic environment was limited and unsubstantiated; and a high-order language tailored specifically for aerospace applications was nonexistent. Even at this early juncture, however,

NASA had begun developing a high-order software language—HAL/S—for the shuttle. This software would ultimately become the standard for Orbiter operations during the Space Shuttle Program.

Software Capability Beyond Technology Limits

NASA contemplated the number of necessary computer configurations during the early stages of Space Shuttle development. It took into consideration the segregation of flight control from guidance and navigation, as well as the relegation of mechanized aerodynamic ascent/re-entry and spaceflight functions to different machines.

These considerations led to a tightly coupled, synchronized fail-operational/fail-safe computation requirement for flight control and sequencing functions that drove the system toward a four-machine computer complex. In addition, the difficulties NASA faced in attempting to interconnect and operate multiple complexes of machines led to the development of a single complex with central integrated computation.

NASA added a fifth machine for off-loading nonessential mission applications, payload, and system-management tasks from the other four machines. Although this fifth computer was also positioned to handle the additional computation requirements that might be placed on the system, it eventually hosted the backup system flight software.

The space agency had to determine the size of the Orbiter computer memory to be baselined and do so within the constraints of computer design and vehicle structure. Memory limitations posed a formidable

Personal Computer Ground Operations Aerospace Language Offered Engineers a “View”

Personal Computer Ground Operations Aerospace Language (PCGOAL) was a custom, PC-based, certified advisory system that provided engineers with real-time data display and plotting. The enhanced situational awareness aided engineers with the decision-making process and troubleshooting during test, launch, and landing operations.

When shuttle landings first began at Dryden Flight Research Center (DFRC), California, Kennedy Space Center (KSC) engineers had limited data-visualization capability. The original disk operating system (DOS)-based PCGOAL first supported KSC engineers during the STS-34 (1989) landing at DFRC. Data were sent from KSC via telephone modem and engineers had visibility to the Orbiter data on site at DFRC. Firing room console-like displays provided engineers with a familiar look of the command and control displays used for shuttle processing and launch countdown, and the application offered the first high-resolution, real-time plotting capability.

PCGOAL evolved with additional capabilities. After design certification review in 1995, the application was considered acceptable for decision making in conjunction with the command and control applications in the firing rooms and DFRC. In 2004, the application was given a new platform to run on a Windows 2000 operating system.

As the Windows-based version of PCGOAL was being deployed, work had already begun to add visualization capabilities. The upgraded application and upgraded editor were deployed in December 2005 at KSC first and later at DFRC and Marshall Space Flight Center/ Huntsville Operations Support Center.



challenge for NASA early in the development phase; however, with the technological advancements that soon followed came the ability to increase the amount of memory.

NASA faced much skepticism from within its organization, regarding the viability of using a high-order language. Assembly language could be used to produce compact, efficient, and fast software code, but it was very similar in complexity to the computer's machine language and therefore required the programmer to understand the intricacies of the computer hardware and instruction set. For example, assembly language addressed the machine's registers directly and operations on the data in the registers directly.

While it might not result in as fast and efficient a code, using a high-order programming language would provide abstraction from the details of the computer hardware, be less cryptic and closer to natural language, and therefore be easier to develop and maintain. As the space agency contracted for the development of HAL/S, program participants questioned the software's ability to produce code with the size, efficiency, and speed comparable to those of an assembly language program. All participants, however, supported a top-down structured approach to software design.

To resolve the issue and quell any fears as to the capability of HAL/S, NASA tested both options and discovered that the nominal loss in efficiency of the high-order language was insignificant when compared to the advantages of increased programmer productivity, program maintainability, and visibility into the software. Therefore, NASA selected HAL/S for all but one software module (i.e., operating system software), thus fulfilling the remaining baselined requirements and approach.

Operating Software for Avionics System

The Orbiter avionics system operation required two independent software systems with a distinct hierarchy and clear delegation of responsibilities. The Primary Avionics Software System was the workhorse of the two systems. It consisted of several memory loads and performed mission and system functions. The Backup Flight System software was just that: a backup. Yet, it played a critical role in the safety and function of the Orbiter. The Backup Flight System software was composed of one memory load and worked only during critical mission phases to provide an alternate means of orbital insertion or return to Earth in the event of a Primary Avionics Software System failure.

Primary Avionics Software System

The Primary Avionics Software System performed three major functions: guidance, navigation, and control of the vehicle during flight; the systems management involved in monitoring and controlling vehicle subsystems; and payload—later changed to vehicle utility—involving preflight checkout functions.

The depth and complexity of Orbiter requirements demanded more memory capacity than was available from a general purpose computer. As a solution, NASA structured each of the major functions into a collection of programs and capabilities needed to conduct a mission phase or perform an integrated function. These collections were called “operational sequences,” and they formed memory configurations that were loaded into the general purpose computers from on-board tape units. Memory overlays were inevitable; however, to a great extent NASA structured these overlays only in quiescent, non-dynamic periods.

The substructure within operational sequences was a choreographed network consisting of major modes, specialist functions, and display functions. Major modes were substructured into blocks that segmented the processes into steps or sequences. These blocks were linked to cathode ray tube display pages so the crew could monitor and control the function. The crew could initiate sequencing through keyboard entry. In certain instances, sequencing could be initiated automatically by the software. Blocks within the specialist functions, initiated by keyboard entry, were linked to cathode ray tube pages. These blocks established and presented valid keyboard entry options available to the crew for controlling the operation or monitoring the process. Major modes accomplished the primary functions within a sequence, and specialist functions were used for secondary or background functions. The display functions, also initiated by keyboard input, contained processing necessary to produce the display and were used only for monitoring data processing results.

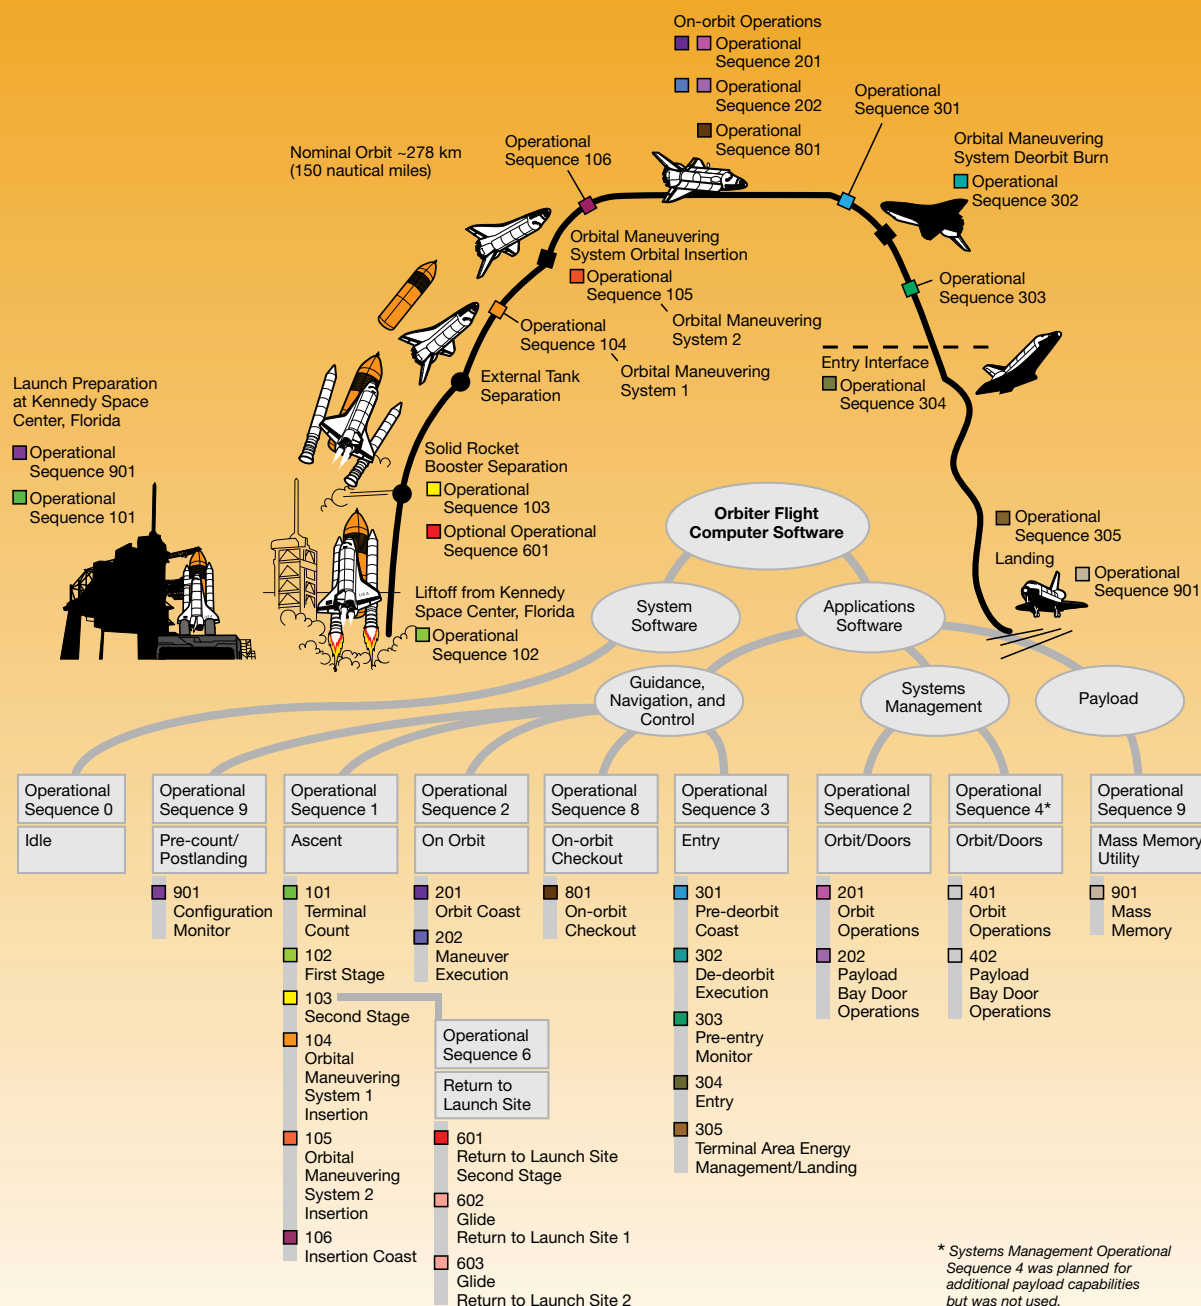
Backup Flight System

The Backup Flight System remained poised to take over primary control in the event of Primary Avionics Software System failure, and NASA thoroughly prepared the backup system for this potential problem. The system consisted of the designated general purpose computer, three backup flight controllers, the backup software, and associated switches and displays.

As far as designating a specific general purpose computer, NASA did not favor any particular one over the others—any of the five could be designated the backup machine by appropriate keyboard entry. The designated computer would request the backup



Mission Phase With Corresponding Operational Sequences and Major Modes



Due to computer memory limitations, the flight software was divided into a number of separate programs called operational sequences. Each sequence provided functions specific to a particular mission phase and were only loaded into memory during that phase of flight.



The Ground Launch Sequencer Orchestrated Launch Success

During launch countdown, the ground launch sequencer was like an orchestra's conductor. Developed in 1978, the sequencer was the software supervisor of critical command sequencing and measurement verification from 2 hours before launch time to launch time and through safing, thus assuring a steady and an appropriate tempo for a safe and successful launch.

Engineered to expedite and automate operations and maximize automatic error detection and recovery, the ground launch sequencer focused on "go/no-go" criteria. Responding to a no-go detection, it could initiate a countdown hold, abort, or

recycle or contingency operations. While controlling certain monitoring aspects, the sequencer did not reduce the engineer's capability to monitor his or her system's health/integrity; however, by assuming command responsibility, it integrated launch requirements and activities, and reduced communication traffic and required hardware. Manual intervention was available for off-nominal conditions.

The four ground launch sequencer components included: exception monitoring; sequencer; countdown clock control; and safing. For exception monitoring, the sequencer continuously monitored more than 1,200 measurements.

If a measurement violated its expected value, the sequencer checked whether the measurement was part of a voting logic group. If voting failed, it automatically caused the countdown to hold at the next milestone or abort the countdown.

The sequencer provided a single point of control during countdown, issuing all commands to ground and flight equipment from the designated period called T minus 9 minutes (T=time) through liftoff. It verified events required for liftoff. If an event wasn't completed, an automated hold/recycle was requested.

Clock control provided the required synchronization between ground and vehicle systems and managed countdown holds/recycles. Clock control allowed the sequencer to resume the countdown after a problem was resolved. The safing component halted the Orbiter's on-board software and, based on the progression of the sequencer, commanded ground and flight systems into a safe configuration for crew egress.



Launch countdown operations in Firing Room 4 at Kennedy Space Center, Florida.

software load from mass memory. The backup computer would then remain on standby. During normal operations, when the primary system controlled the Orbiter, the backup system operated in "listen" mode to monitor and obtain data from all prime machines and their assigned sensors. By acquiring these data, the Backup Flight System maintained computational currency and, thus, the capability to assume control of the Orbiter at any time.

NASA independently developed and coded the software package for the

Backup Flight System as an added level of protection to reduce the possibility of generic software errors common to the primary system. The entire Backup Flight System was contained in one memory configuration, loaded before liftoff, and normally maintained in that machine.

Success—On Multiple Levels

NASA overcame the obstacles it faced in creating the shuttle's Primary Avionics Software System through ingenuity and expertise. Even

technology that was current during the initial planning stages did not impose limits on what the space agency could accomplish in this area. NASA succeeded in pushing the boundaries for what was possible by structuring a system that could handle multiple functions within very real parameters. It also structured a backup support system capable of handling the demands of spaceflight at a critical moment's notice.



Integrated Extravehicular Activity/Robotics Virtual Reality Simulation

As the Space Shuttle Program progressed into the 1990s, the integration of extravehicular activity (EVA) and robotics took on a whole new importance when Hubble Space Telescope servicing/repair (first flight 1993) and space-based assembly of the International Space Station (ISS) tasks were realistically evaluated.

Two motivating factors influenced NASA's investigation into the potential use of virtual reality technology that was barely in its infancy at that time. The first factor was in response to a concern that once Hubble was deployed on orbit future astronauts and flight controllers would not have easy access to the telescope to familiarize themselves with the actual hardware configuration to plan, develop, and review servicing procedures.

The second factor was based on previous on-orbit experience with the interaction and communication between EVA crew members and Shuttle Robotic Arm operators. NASA discovered that interpreting instructions given by a crew member located in a foot restraint on the end of the robotic arm was not as intuitive to the arm operator as first thought, especially when both were not in the same body orientation when giving or receiving commands. The EVA crew member could, for example, be upside down with respect to the robotic arm operator in microgravity. Therefore, the command to "Move me up" left the arm operator in a quandary trying to decide what "up" actually meant.

NASA Embraces Advances in Virtual Reality

It was at this same time in the early 1990s that virtual reality hardware started to enter the commercial world in the form of head-mounted displays, data gloves, motion-tracking instruments, etc.

In the astronaut training world, no facility allowed an EVA crew member to ride on a robotic arm operated by another crew member in a realistic space environment. The Water Emersion Test Facility at Johnson Space Center (JSC) in Houston, Texas, provided a training arena for EVA crew members, but the confined space and the desire to not require subjects to be heads down for more than very short periods of time did not allow for suitable integrated training between the EVA crew and the robotic arm operators. Likewise, the Manipulator Development Facility's hydraulic arm and the computer graphic-based robotic arm simulators at JSC were not conducive for EVA crew interaction.

Virtual reality provided a forum to actually tie those two training scenarios together in one simulation. Working closely with the astronaut office, NASA

engineers took commercially available virtual reality hardware and developed the computer graphic display software and across-platform communications software that linked into existing "man-in-the-loop" robotic arm computer simulations to produce an integrated EVA/robotics training capability.

Virtual Reality Is Put to the Test

The first use of these new capabilities was in support of crew training for Space Transportation System (STS)-61 (1993)—the Hubble Space Telescope servicing mission. The virtual reality simulation provided a flight-like environment in which the crew was able to develop and practice the intricate choreography between the Shuttle Robotic Arm operator and the EVA crew member affixed to the end of that arm. The view in the head-mounted display was as it would be seen by the astronaut working around the Hubble berthed in the shuttle payload bay at an orbital altitude of 531 km (330 miles) above the Earth.

The next opportunity to take advantage of the virtual reality software involved EVA crew members training to perform the first engineering test flights of the



Astronaut Mark Lee trains for his Simplified Aid for EVA Rescue test flight (STS-64 [1994]) using the virtual reality flight trainer (left) and on orbit (right).



Simplified Aid for EVA Rescue (SAFER) on STS-64 (1994).

The output of a dynamic simulation of the SAFER backpack control system and its flying characteristics, using zero-gravity as a parameter, drove the head-mounted display visual graphics. Inputs to the simulation were made using a flight-equivalent engineering unit hand controller. The EVA crew member practiced and refined the flight test maneuvers to be flown during on-orbit tests of the rescue unit. The crew member could see the on-orbit configuration of the shuttle payload bay, the robotic arm, and the Earth/horizon through the virtual reality head-mounted display at the orbital altitude planned for the mission. The EVA crew member was also able to interact with the robotic arm operator as well as see the motions of the arm, which was an integral part of the on-orbit tests. The robotic arm operator was also able to view the EVA crew member's motions in the simulated shuttle payload bay camera views made available to the operator as part of the dynamic man-in-the-loop robotic arm simulation.

As a result of the engineering flights of the SAFER unit on STS-64, NASA was able to validate the virtual reality simulation and it became the ground-based SAFER training simulator used by all EVA crew members assigned to space station assembly missions.


Each EVA crew member was required to have at least four 2-hour training classes prior to a flight to practice flying rescue scenarios with the unit in the event he or she became separated from the space vehicle during an EVA.

NASA also developed a trainer that was flown on board the space station laptop computers. The trainer used the same simulation and display software as the ground-based simulator, but it incorporated a flat-screen display instead of a head-mounted display. It also used the same graphic model database as the ground-based simulators. ISS crew members used the on-board trainer to maintain SAFER hand controller proficiency throughout their time on the ISS.


Handling Large Objects During Extravehicular Activity

Learning to handle large objects in the weightlessness of space also posed a unique problem for EVA crew members training in ground-based facilities. In the microgravity environment of space, objects may be weightless but they still have mass and inertia as well as a mass distribution around a center of gravity.

NASA engineers developed a tendon-driven robot and a set of dynamic control software to simulate the feel and motion of large objects being handled by an EVA crew member within the zero-gravity parameter. The basic concept was to mount a reel of cable and an electric drive motor at each of the eight corners of a structure that measured approximately 3 m (10 ft) on a side. Each cable was then attached to one of the eight corners of an approximately 0.6-m (2-ft) cube. In this configuration, the position and orientation of the smaller cube within the large structure could be controlled by reeling in and out the cables. Load cells were mounted to the smaller cube



Astronauts Richard Linnehan (above left) and Nancy Currie (below) use the zero-gravity mass handling simulation and the Shuttle Robotic Arm simulation to practice combined operations prior to flight. The large image on the right is a rendering of the simulation. The inset is an actual photo of Astronaut Richard Linnehan (STS-109 [2002]) unfolding a solar array while anchored to the end of the robotic arm.



International Space Station Expedition 10 crew members Leroy Chiao (left) and Salizhan Sharipov train in virtual reality to photograph an approaching Orbiter through the space station windows. The lower pictures show what each sees through his respective camera view finder.

while handrails or other handling devices were attached to the load cells. As a crew member applied force to the handling device, the load cells measured the force and fed those values to a dynamic simulation that had the mass characteristics of the object being handled as though it were in weightlessness. Output from the computer program then drove the eight motors to move the smaller cube accordingly. Once these elements were integrated into graphics in the head-mounted display, the crew member not only felt the resulting six-degree-of-freedom motion of the simulated object, he or she also saw a three-dimensional (3-D) graphical representation of the real-world object in its actual surrounding environment.

The mass handling simulation—called kinesthetic application of mechanical force reflection—was qualitatively validated over a number of shuttle flights starting with STS-63 (1995). On that flight, EVA crew members were scheduled to handle objects that weighed from 318 to 1,361 kg (700 to 3,000 pounds) during an EVA. After their flight, they evaluated the ability of the application to simulate the handling conditions experienced in microgravity.

Kinesthetic application of mechanical force reflection was deemed able to faithfully produce an accurate simulation of the feel of large objects being handled by EVA crew members following a number of postflight evaluations.

Kinesthetic application of mechanical force reflection was also integrated with the Shuttle Robotic Arm simulation, which allowed the EVA crew member riding on the end of the arm to actually feel the arm-induced motion in a large payload that he or she would be holding during a construction or repair operation around the ISS or Hubble.

NASA built two kinesthetic application of mechanical force reflections so that two EVA crew members could train to handle the same large object from two different vantage points. The forces and motion input by one crew member were felt and seen by the other crew member. This capability allowed crew members to evaluate mass handling techniques preflight. It also allowed them to work out not only the command protocol they planned to use, but also which crew member would be controlling the object and which would be stabilizing the object during the EVA.

Virtual Reality Simulates On-orbit Conditions

Following the Columbia accident in 2003, as a shuttle approached the space station, space station crew members photographed its Thermal Protection System from a distance of 183 m (600 ft) using digital cameras with 400mm and 800mm telephoto lenses.

As in previous scenarios, there was no place on Earth where crew members could practice photographing a Space Shuttle doing a 360-degree pitch maneuver at a distance of 183 m (600 ft). Virtual reality was again used to realistically simulate the on-orbit conditions and provide ground-based training to all space station crew members prior to their extended stay in space.

Engineers placed a cathode ray tube display from a head-mounted display inside a mocked-up telephoto lens. The same 3-D graphic simulation that was used to support the previous applications drove the display in the telephoto lens to show a shuttle doing the pitch maneuver at a range of 183 m (600 ft). With a real camera body attached to the mocked-up lens, each crew member could practice photographing the shuttle during its approach maneuver.

Summary

NASA took advantage of the benefits that virtual reality had to offer. Beginning in 1992, the space agency used the technology at JSC to support integrated EVA/robotics training for all subsequent EVA flights, including SAFER engineering flights, Hubble repair/servicing missions, and the assembly and maintenance of the ISS. Each EVA crew member spent from 80 to 120 hours using virtual reality to train for work in space.



Integrated Solutions for Space Shuttle Management... and Future Endeavors

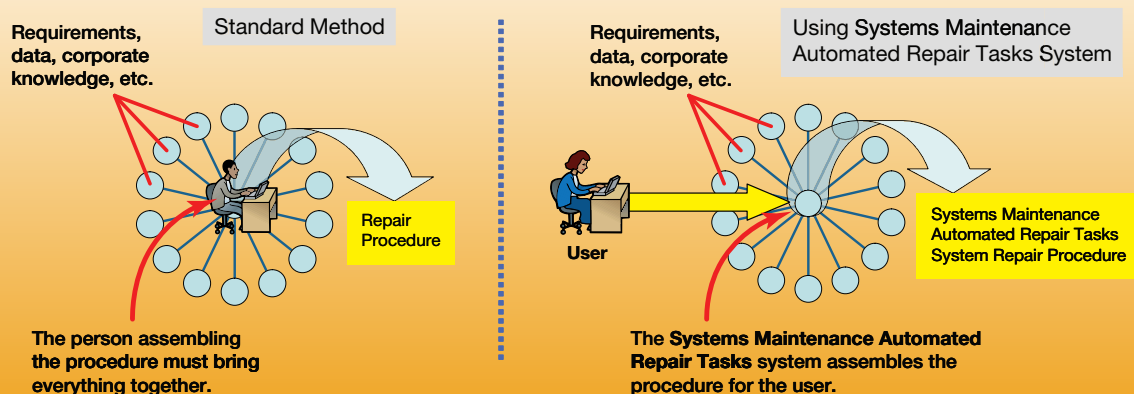
Kennedy Space Center (KSC) developed an integrated, wireless, and paperless computer-based system for management of the Space Shuttle and future space program products and processes. This capability was called Collaborative Integrated Processing Solutions. It used commercial off-the-shelf software products to provide an end-to-end integrated solution for requirements management, configuration management, supply chain planning, asset life cycle management, process engineering/process execution, and integrated data management. This system was accessible from stationary workstations and tablet computers using wireless networks.

Collaborative Integrated Processing Solutions leveraged the successful implementation of Solumina® (iBASEt, Foothill Ranch, California)—a manufacturing execution system that provided work instruction authorization, electronic approval, and paperless work execution. Solumina® provided real-time status updates to all users working on the same document. The system provided for

electronic buy off of work instructions, electronic data collection, and embedded links to reference materials. The application included electronic change tracking and configuration management of work instructions. Automated controls provided constraints management, data validation, configuration, and reporting of consumption of parts and materials.

In addition, KSC developed an interactive decision analysis and refinement software system known as Systems Maintenance Automated Repair Tasks. This system used evaluation criteria for discrepant conditions to automatically populate a document/procedure with predefined steps for safe, effective, and efficient repair. It stored tacit (corporate) knowledge, merging hardware specification requirements with actual “how-to” repair methods, sequences, and required equipment. Although the system was developed for Space Shuttle applications, its interface is easily adaptable to any hardware that can be broken down by component, subcomponent, discrepancy, and repair.

Systems Maintenance Automated Repair Tasks Solution Philosophy—Variables



The Systems Maintenance Automated Repair Tasks allowed corporate knowledge to be kept in-house while increasing efficiency and lowering cost.

Three-Dimensional Graphics Provide Extraordinary Vantage Points

Astronauts' accomplishments in space seem effortless, yet they spent many hours on the ground training and preparing for missions.

Some of the earliest engineering concept development and training took place in the Johnson Space Center

Virtual Reality Laboratory and involved the Dynamic Onboard Ubiquitous Graphics (DOUG) software package. NASA developed this three-dimensional (3-D) graphics-rendering package to support integrated training among the Shuttle Robotic Arm operators, the International Space Station (ISS) Robotic Arm operators, and the extravehicular activity (EVA) crew members. The package provided complete software and model database commonality among ground-based crew training simulators, ground-based

EVA planning tools, on-board robotic situational awareness tools, on-board training simulations, and on-board EVA/robotic operations review tools for both Space Shuttle and ISS crews.

Level-of-detail Capability


Originally, the software was written as an application programming interface—an interface that enables the software to interact with other software—around the graphics-rendering package developed to support the virtual reality

Additional Extravehicular Activity Support



The International Space Station (ISS) has more than 2,300 handrails located on its exterior. These handrails provide translation paths for extravehicular activity (EVA) crew members. Pull-down menus in the Dynamic Onboard Ubiquitous Graphics (DOUG) software allow the user to highlight and locate each handrail. Entire translation paths can be highlighted and displayed for review by crew members prior to performing an EVA.

More than 620 work interface sockets are located on the external structure of the ISS, and nine articulating portable foot restraints can be relocated to any of the work interface sockets. Each articulating portable foot restraint has three articulating joints and a rotating base that produce 33,264 different orientations for an EVA crew member standing in that particular foot restraint. Each work interface socket can be located in the software package, and each articulating portable foot restraint can be configured to show all potential worksites and worksite configurations to support EVA planning.

The DOUG software package also contains and can highlight the locations of externally mounted orbital replacement units on the ISS, thruster and antenna keep-out zones that affect EVA crew member positioning, and articulating antennas, radiators, and solar arrays—all of which are configurable.



Articulated portable foot restraints configuration (top) and highlighted translation path (bottom).



These two views show the effect of level-of-detail control. The left view is a high-resolution image compared to the low-resolution image on the right.

training simulation. The Simplified Aid for EVA Rescue (SAFER) on-board trainer required software that would run on the original IBM 760 laptop computers on board the ISS and thus required the UNIX-based code to be ported to a Windows-based operating system. The limited graphics capability of those computers also required additional model database artifacts that provided level-of-detail manipulation to make the simulation adequate for its intended purpose. This additional level-of-detail capability allowed the same high-fidelity model database developed for EVA training in the virtual reality facility to be used on the laptop computers on the ISS.

To obtain adequate graphics performance and screen update rates for simulating SAFER flying, crew members could select a low level-of-detail scene, which still displayed enough detail for the recognition of station landmarks and motion cues.

The DOUG software package, when not in use as a trainer, also provided a highly detailed, interactive 3-D model of the ISS that was viewable from any vantage point via keyboard inputs. The software first flew on board both shuttle and station in March 2001, and during


Space Transportation System (STS)-102, and was on all subsequent shuttle and station flights with the exception of STS-107 (2003). That flight did not carry a robotic arm, had no planned EVAs, and did not dock with the ISS.

Benefits for Robotic Arm Operations

The DOUG software package supported SAFER training. The software was also capable of providing the situational

awareness function during Space Station Robotic Arm operations by connecting to the on-board payload general support computer and using the telemetry from the arm to update the graphic representation in the program display.

The same software was compatible with laptop computers flown on the shuttle, and the graphical Shuttle Robotic Arm could be similarly driven with shuttle arm telemetry. Different viewpoints




Dynamic Onboard Ubiquitous Graphics displays multiple simulated camera and synthetic eye-point views on the same screen. The simulated camera views show the Japanese Experiment Module and the Columbus Laboratory in the top left image, the Mini Research Module-1 in the top right image, and the International Space Station in the bottom image.

could be defined in the software to represent the locations of various television cameras located around station and shuttle. The various camera parameters were defined in the software to display the actual field of view, based on the pan and tilt capabilities as well as the zoom characteristics of each camera.

The second ISS crew (2001) used these initial capabilities to practice for upcoming station assembly tasks with the Space Station Robotic Arm prior to the actual components arriving on a shuttle flight. The crew accomplished this by operating the real robotic arm using the real hand controllers and configuring a “DOUG laptop” to receive remote manipulator joint angle telemetry.

The graphics contained the station configuration with the shuttle docked and the station airlock component located in the shuttle’s payload bay. The arm operator could see synthetic end-effector camera views produced in the program. These views showed the airlock with its grapple fixture in the payload bay of the Orbiter even though no Orbiter actually existed. The operator practiced maneuvering the real arm end-effector onto an imaginary grapple fixture and then maneuvering the real arm with the imaginary airlock attached, through the prescribed trajectory to berth the imaginary airlock onto the real common berthing mechanism on the ISS Unity Node.

Through DOUG the arm operator also had access to synthetic views from all the shuttle cameras, as well as the Space Station Robotic Arm cameras that would be used during the actual assembly operations. This made training much more effective than simply driving the robotic arm around in open space.



The colors displayed in Dynamic Onboard Ubiquitous Graphics indicate direction of approach of the robotic arm booms with respect to the closest object: green = opening; yellow = closing; and red = envelope violation.

Proximity Detection

As the ISS grew in complexity, NASA added capabilities to the DOUG software. Following a near collision between the Space Station Robotic Arm and one of the antennas located on the laboratory module of the ISS, the space agency added the ability to detect objects close to one another—i.e., proximity detection. The software calculated and displayed the point of closest approach for the main robotic arm booms and the elbow joint to any station or shuttle component displayed in the model database.

A vector was drawn between each of the three robotic arm components and the nearest structure. When DOUG received robotic arm telemetry data and was being used for situational awareness during robotic arm operations, the color of these vectors indicated whether measured distance was increasing or decreasing. It also indicated whether the relative distance was within a user-defined, keep-out envelope around the robotic arm. Both audible and graphical warnings were selectable to indicate when a keep-out envelope was breached.




Thermal Protection System Evaluation

During the preparation for Return to Flight following the Columbia accident in 2003, NASA incorporated the entire shuttle Thermal Protection

System database and a “painting” feature into the DOUG software package. The database consisted of all 25,000+ tiles, thermal blankets, reinforced carbon-carbon wing leading edge panels, and nose cap.

The software was used preflight to develop the trajectories of the Shuttle Robotic Arm and Orbiter Boom Sensor System used to perform in-flight Orbiter inspections. The software allowed engineers to “paint” the areas that were within the specifications



An example of the tile highlighting and painting feature in Dynamic Onboard Ubiquitous Graphics.



of various sensors on the Orbiter Boom Sensor System (e.g., range, field of view, incidence angle) to make sure the Thermal Protection System was completely covered during on-orbit surveys.

The same configuration models and tile database used on the ground were also loaded on the on-board laptop computers. This allowed the areas of interest found during the survey data analysis to be highlighted and uplinked to the shuttle and station crews for further review using the DOUG program.

Inspection of the STS-114 (2005) survey data showed protruding gap fillers between tiles on the Orbiter. These protrusions were of concern for re-entry into Earth's atmosphere. Ground controllers were able to highlight the surrounding tiles in the database, develop a Space Station Robotic Arm configuration with an EVA crew member in a foot restraint on the end, and uplink that configuration file to the station laptop computers. The crew members were then able to use the software to view the area of concern, understand how they would need to be positioned underneath the Orbiter, get a feel for the types of clearances they had with the structure around the robotic arm, and evaluate camera views that would be available during the operation.

Having the 3-D, interactive viewing capability allowed crew members to become comfortable with their understanding of the procedure in much less time than would have been required with just "words" from ground control. A key aspect to the success of this scenario was the software and

configuration database commonality that DOUG provided to all participants—station and shuttle crews, ground analysis groups, procedure developers, mission controllers, and simulation facilities.

DOUG was loaded on more than 1,500 machines following the Columbia accident and was used as a tool to support preflight planning and procedures development as well as on-orbit reviews of all robotic and EVA operations. In addition to its basic capabilities, the software possessed many other features that made it a powerful planning and visualization tool.

Expansion of Capabilities

DOUG has also been repackaged into a more user-friendly application referred to as Engineering DOUG Graphics for Exploration (EDGE). This application is a collection of utilities, documentation, development tools, and visualization tools wrapped around the original renderer. DOUG is basically the kernel of the repackaged version, which includes the addition of various plug-ins, models, scripts, simulation interface code, graphical user interface add-ons, overlays, and development interfaces to create a visualization package. The project allows groups to quickly visualize their simulations in 3-D and provides common visuals for future program cockpits and training facilities. It also allows customers to expand the capabilities of the original software package while being able to leverage off the development and commonality achieved by that software in the Space Shuttle and ISS Programs.

Summary

The graphics-rendering software developed by NASA to support astronaut training and engineering simulation visualization during the shuttle era provided the cornerstone for commonality among ground-based training facilities for both the Space Shuttle and the ISS. The software has evolved over the years to take advantage of ever-advancing computer graphics technology to keep NASA training simulators state of the art and to provide a valuable resource for future programs and missions.



Structural Design

Introduction

Gail Chapline

Orbiter Structural Design

Thomas Moser

Glenn Miller

Shuttle Wing Loads—Testing and Modification Led to Greater Capacity

Tom Modlin

Innovative Concept for Jackscrews Prevented Catastrophic Failures

John Fraley

Richard Ring

Charles Stevenson

Ivan Velez

Orbiter Structure Qualification

Thomas Moser

Glenn Miller

Space Shuttle Pogo—NASA Eliminates “Bad Vibrations”

Tom Modlin

Pressure Vessel Experience

Scott Forth

Glenn Ecord

Willard Castner

Nozzle Flexible Bearing—Steering the Reusable Solid Rocket Motor

Coy Jordan

Fracture Control Technology Innovations—From the Space Shuttle Program to Worldwide Use

Joachim Beek

Royce Forman

Glenn Ecord

Willard Castner

Gwyn Faile

Space Shuttle Main Engine Fracture Control

Gregory Swanson

Katherine Van Hooser

The Space Shuttle—a mostly reusable, human-rated launch vehicle, spacecraft, space habitat, laboratory, re-entry vehicle, and aircraft—was an unprecedented structural engineering challenge. The design had to meet several demands, which resulted in innovative solutions. The vehicle needed to be highly reliable for environments that could not be simulated on Earth or fully modeled analytically for combined mechanical and thermal loads. It had to accommodate payloads that were not defined or characterized. It needed to be weight efficient by employing a greater use of advanced composite materials, and it had to rely on fracture mechanics for design with acceptable life requirements. It also had to be certified to meet strength and life requirements by innovative methods. During the Space Shuttle Program, many such structural design innovations were developed and extended to vehicle processing from flight to flight.

Orbiter Structural Design

NASA faced several challenges in the structural design of the Orbiter. These challenges were greater than those of any previous aircraft, launch vehicle, or spacecraft, and the Orbiter was all three. Yet, the space agency proceeded with tenacity and confidence, and ultimately reached its goals. In fact, 30 years of successful shuttle flights validated the agency's unique and innovative approaches, processes, and decisions regarding characteristics of design.

A few of the more significant challenges NASA faced in Orbiter structural design included the evolution of design loads. The Orbiter structure was designed to an early set of loads and conditions and certified to a later set. The shuttle achieved first-flight readiness through a series of localized structural modifications and operational flight constraints. During the early design phase, computer analyses using complex calculations like finite-element models and techniques for combined thermal and mechanical loads were not possible. Later advances in analytical methods, coupled with test data, allowed significant reductions in both scope and cost of Orbiter structural certification. The space agency had to face other challenges. Structural efficiency had to be compromised to assure versatile payload attachment and payload bay door operations. Skin buckling had to be avoided to assure compatibility with the low-strength Thermal Protection System tiles. Composite materials

beyond the state of the art were needed. The crew compartment had to be placed into the airframe such that the pressurized volume would effectively "float." And it was impractical to test the full airframe under combined mechanical and thermal loads.

Thousands of analytical design loads and conditions were proven acceptable with flight data with one exception:

the ascent wing loads were greater than predicted because of the effect the rocket exhaust plume had on the aerodynamic pressure distribution. As a result, early flights were flown within limited flight regimes to assure that the structural capability of the wings was not exceeded. The wings were later "strengthened" with minor changes in the design and weight.

Shuttle Wing Loads—Testing and Modification Led to Greater Capacity

Orbiter wing loads demonstrated the importance of anchoring the prediction or grounding the analysis with flight data in assuring a successful flight. The right wing of Columbia was instrumented with strain gauges for the test flights and was load-calibrated to verify the in-flight air load distribution. The wing was also instrumented with pressure gauges; however, the number was limited due to on-board recorder space limitations. This resulted in the need to obtain additional pressure data.

Space Transportation System (STS)-1 (1981) data indicated higher shear in the aft spar web than was predicted. NASA conducted analyses to determine the location and magnitude of forces causing this condition. The results indicated an additional load along the outboard wing leading edge (elevon hinge line). Data obtained on STS-2 (1981) through STS-4 (1982) substantiated these results. This caused concern for the operational wing limits that were to be imposed after the flight test period.

The additional load caused higher bending and torsion on the wing structure, exceeding design limits. The flight limits, in terms of angle of attack and sideslip, would have to be restricted with an attendant reduction in performance.

The recovery plan resulted in modification to the wing leading edge fittings. The major impact was to the structure between the upper and lower wing skins, which were graphite-epoxy. These required angle stiffeners on each flat to increase the buckling stress. The weight of the modifications resulted in a loss of performance. The resulting flight envelope was slightly larger than the original when accounting for the negative angle-of-attack region of the flight regime.


Payload Access and Structural Attachments—Mid-Fuselage and Payload Bay Doors

NASA designed the mid-fuselage of the Orbiter to be “flexible” so as to accommodate the closing of payload bay doors in space. The design also had to accommodate a wide range of payload sizes, weights, and number.

The payload bay doors were an integral part of the fuselage structure. The classical structural design would have the doors provide strength when the fuselage encountered loads from bending, twisting, shear, internal pressure, and thermal gradients. The doors also had to open in space to provide access to the payload and enable the radiators to radiate heat to space. Equally important, the doors had to close prior to re-entry into Earth’s atmosphere to provide aerodynamic shape and thermal protection.

To balance the functional and strength requirements, engineers designed the doors to be flexible. The flexibility and zipper-like closing ensured that the doors would close in orbit even if distorted thermally or by changes in the gravity environment (from Earth gravity to microgravity). If the latches did not fully engage, the doors could not be relied on to provide strength during re-entry for fuselage bending, torsion, and aerodynamic pressure. Thus, the classical design approach for ascent was not possible for re-entry. The bulkheads at each end of the payload section and the longerons on each side required additional strength. To reduce weight and thermal distortion, engineers designed the doors using graphite epoxy. This was the largest composite structure on any aircraft or spacecraft at the time.

Typical Payload Attachment Scheme



Sets of moveable attachment fittings on the longerons and frames accommodated multiple payloads. The Monte Carlo analyses of the full spectrum of payload quantities, sizes, mass properties, and locations determined the mid-fuselage design loads. These design loads were enveloped based on a combination of 10 million load cases. Decoupling the design of the mid-fuselage and payloads enabled a timely design of both.

The mid-fuselage had to accommodate the quantity, size, weight, location, stiffness, and limitations of known and unknown payloads. An innovative design approach needed to provide a statically determinant attachment system between the payloads and mid-fuselage. This would decouple the bending, twisting, and shear loads between the two structures, thus enabling engineers to design both without knowing the stiffness characteristic of each.

Designing to Minimize Local Deflections

The Orbiter skin was covered with more than 30,000 silica tiles to withstand the heat of re-entry. These tiles had a limited capacity to accommodate structural deflections from thermal gradients. The European supersonic Concorde passenger aircraft (first flown in 1969 and in service from 1976 to 2003) and the SR-71 US military

aircraft encountered significant thermal gradients during flight. The design approach in each was to reduce stresses induced by the thermal gradients by enabling expansion of selected regions of the structure; e.g., corrugated wing skins for the SR-71 and “slots” in the Concorde fuselage. After consulting with the designers of both aircraft, NASA concluded that the Orbiter design should account for thermally induced stresses but resist large expansions and associated skin buckling. This brute-force approach

protected the attached silica tile as well as simplified the design and manufacture of the Orbiter airframe.

NASA developed these design criteria so that if the thermal stresses reduced the mechanical stresses, the reductions would not be considered in the combined stress calculations.

To determine the thermally induced stresses, NASA established deterministic temperatures for eight initial temperature conditions on the Orbiter at the time of re-entry as well

as at several times during re-entry. Engineers generated 120 thermal math models for specific regions of the Orbiter. Temperatures were extrapolated and interpolated to nodes within these thermal math models.


Use of Unique Advanced Materials

Even though the Orbiter was a unique aircraft and spacecraft, NASA selected a conventional aircraft skin/stringer/frame design approach. The space agency also used conventional aircraft material (i.e., aluminum) for the primary structure, with exceptions in selected regions where the use of advanced state-of-the-art composites increased efficiency due to their lower density, minimum thermal expansion, or higher modulus of elasticity.

Other exceptions to the highly reliable conventional structures were the graphite-epoxy Orbital Maneuvering System skins, which were part of a honeycomb sandwich structure. These graphite honeycomb structures had a vented core to relieve pressure differentials across the face sheets during flight. They also required a humidity-controlled environment while on the ground to prevent moisture buildup in the core. Such a buildup could become a source of steam during the higher temperature regimes of flight. Finally, during the weight-savings program instituted on Discovery, Atlantis, and Endeavour, engineers replaced the aluminum spar webs in the wing with a graphite/epoxy laminate.

Large doors, located on the bottom of the Orbiter, were made out of beryllium. These doors closed over the External Tank umbilical cavity once the vehicle


Orbiter Thermal Stress Analysis Modeling



Early Trade Studies Showed Cost Benefits That Guided Materials Selection

Titanium offered advantages for the primary structure because of higher temperature capability—315°C vs. 177°C (600°F vs. 350°F). When engineers considered the combined mass of the structure and Thermal Protection System, however, they noted a less than 10% difference. The titanium design cost was 2.5 times greater. The schedule risk was also greater. NASA considered other combinations of materials for the primary structure and Thermal Protection System and conducted a unit cost comparison. This study helped guide the final selections and areas for future development.

Orbiter Structure/Thermal Protection System First Unit Cost Comparison



was on orbit. These approximately 1.3-m (50-in.) square doors maintained the out-of-plane deflection to less than

20 mm (0.8 in.) to avoid contact with adjacent tiles. They also had the ability to withstand a 260°C (500°F)


environment generated by ascent heating. The beryllium material allowed the doors to be relatively lightweight and very stiff, and to perform well at elevated temperatures. The superior thermal performance allowed the door, which measured 25.4 mm (1 in.) in thickness, to fly without internal insulation during launch. Since beryllium can be extremely toxic, special procedures applied to those working in its vicinity.

The truss structure that supported the three Space Shuttle Main Engines was stiff and capable of reacting to over a million pounds of thrust. The 28 members that made up the thrust structure were machined from diffusion-bonded titanium. Titanium strips were placed in an inert environment and bonded together under heat, pressure, and time. This fused the titanium strips into a single, hollow, homogeneous mass. To increase the stiffness, engineers bonded layers of boron/epoxy to the outer surface of the titanium beams. The titanium construction was reinforced in select areas with boron/epoxy tubular struts to minimize weight and add stiffness. Overall, the integrated metallic composite construction reduced the thrust structure weight by 21%, or approximately 409 kg (900 pounds).

NASA used approximately 168 boron aluminum tubes in the mid-fuselage frames as stabilizing elements. Technicians bonded these composite tubes to titanium end fittings and saved approximately 139 kg (305 pounds) over a conventional aluminum tube design. During ground operations, however, composite tubes in high traffic areas were repeatedly damaged and were eventually replaced with an aluminum design to increase robustness during vehicle turnaround.

Orbiter Structure—Structural Arrangement and Location of Composite Materials

Conventional Aluminum Structure
Maximum Temperature 177°C (350°F)
Protected by Reusable Surface Insulation




After the initial design of Challenger and Columbia, NASA initiated a weight-savings program for the follow-on vehicles—Discovery, Atlantis, and Endeavour. The space agency achieved weight savings through optimization of aluminum structures and replaced the aluminum spar webs in the wing with a graphite/epoxy laminate.

“Floating” Crew Compartment

The crew compartment structure “floated” inside the forward fuselage. The crew compartment was attached

to the forward fuselage at four discrete points, thus enabling a simpler design (for pressure and inertia loads only) and greater thermal isolation. The crew compartment was essentially a pressure vessel and the only pressurized compartment in the Orbiter. To help assure pressure integrity, the aluminum design withstood a large noncritical crack while maintaining cabin pressure. The “floating” crew compartment reduced weight over an integrated forward fuselage design and simplified manufacturing.



The crew cabin being installed in the forward fuselage.

© Rockwell International. All rights reserved.

Innovative Concept for Jackscrews Prevented Catastrophic Failures



More than 4,000 jackscrews were in use around Kennedy Space Center (KSC) during the Space Shuttle era. NASA used some of these jackscrews on critical hardware. Thus, a fail-safe, continue-to-operate design was needed to mitigate the possibility of a catastrophic event in case of failure.

A conventional jackscrew contained only one nut made of a material softer than that of the threaded shaft. With prolonged use, the threads in the nut would wear away. If not inspected and replaced after excessive wear, the nut eventually failed. KSC's fail-safe concept for machine jackscrews incorporated a redundant follower nut that would begin to bear the axial jack load on the failure of the primary nut.

Unlike the case of a conventional jackscrew, it was not necessary to relieve the load to measure axial play or disassemble the nut from the threaded shaft to inspect the nut for wear. Instead, wear could be determined by measuring the axial gap between the primary nut and the follower nut.

Additionally, electronic and mechanical wear indicators were used to monitor the gap during operation or assist during inspection. These devices would be designed to generate a warning when the thread was worn to a predetermined thickness. The fail-safe, continue-to-operate design concept offered an alternative for preventing catastrophic failures in jackscrews, which were used widely in aeronautical, aerospace, and industrial applications.

Orbiter Structure Qualification

The conventional strength and life certification approach for a commercial or military aircraft is to demonstrate the ultimate strength and fatigue (life) capacities with a dedicated airframe for each. Similarly, NASA planned two full-scale test articles at the outset of the Orbiter design, development, test, and evaluation program. Ultimately, the Orbiter structure was certified with an airframe that became a flight vehicle and a series of smaller component test articles that comprised about 30% of the flight hardware. The space agency did not take additional risks, and the program costs for ground tests were reduced by several hundred million dollars.

Ultimate Strength Integrity

Virtually all of the Orbiter's primary structure had significant thermal stress components. Therefore, thermal stress had to be accounted for when certifying the design for ultimate strength. Yet, it was impractical—if not impossible—to simulate the correct combination of temperatures and mechanical loads for the numerous conditions associated with ascent, spaceflight, and re-entry into Earth's atmosphere, especially for transient cases of interest. NASA reached this conclusion after consulting with the Concorde aircraft structural experts who conducted multiyear, expensive combined environment tests.

Orbiter strength integrity would be certified in a bold and unconventional approach that used the Challenger (Orbiter) as the structural test article. Rather than testing the ultimate load (140% of maximum expected loads), NASA would test to 120% of limit

mechanical load, use the test data to verify the analytical stress models, and analytically prove that the structure could withstand 140% of the combined mechanical and thermal stresses.

The structural test article was mounted in a horizontal position at the External Tank reaction points and subjected to a ground test program at the Lockheed test facility in Palmdale, California. The 390,900-kg (430-ton) test rig contained 256 hydraulic jacks that distributed loads across 836 application points to simulate various stress levels. Initial influence coefficient tests involved the application of approximately 150 load

conditions as point loads on the vehicle. These unit load cases exercised the structure at the main engine gimbal and actuator attachments, payload fittings, and interfaces on the wing, tail, body flap, and Orbital Maneuvering System pods. Engineers measured load vs. strain at numerous locations and then used those measurements for math model correlation. They also used deflection measurements to substantiate analytical stiffness matrices.

The Orbiter airframe was subjected to a series of static test conditions carried to limit plus load levels (approximately 120% of limit). These conditions

consisted of a matrix of 30 test cases representative of critical phases (boost, re-entry, terminal area energy management, and landing) to simulate design mechanical loads plus six thrust vector-only conditions. These tests verified analytically predicted internal load distributions. In conjunction with analysis, the tests also confirmed the structural integrity of the Orbiter airframe for critical design limit loads. Engineers used these data to support evaluation of the ultimate factor of safety by analysis. Finally, they used the test series to evaluate strains from the developmental flight instrumentation.

Space Shuttle Pogo—NASA Eliminates “Bad Vibrations”

Launch vehicles powered by liquid-fueled, pump-fed rocket engines frequently experience a dynamic instability that caused structural vibrations along the vehicle's longitudinal axis. These vibrations are referred to as “Pogo.”

As Astronaut Michael Collins stated, “The first stage of Titan II vibrated longitudinally so that someone riding on it would be bounced up and down as if on a pogo stick.”

In technical terms, Pogo is a coupled structure/propulsion system instability caused by oscillations in the propellant flow rate that feeds the engines. The propellant flow rate oscillations can result in oscillations in engine thrust. If a frequency band of the thrust oscillations is in phase with the natural frequency of engine structure and is of sufficient magnitude to overcome structural damping, the amplitude of the propellant flow rate oscillation will increase. Subsequently, this event will increase the amplitude of the


thrust oscillation. This sequence can lead to Pogo instability, with the possible result in an unprogrammed engine shutdown and/or structural failure—both of which would result in loss of mission.

Most NASA launch vehicles experienced Pogo problems. Unfortunately, the problem manifested itself in flight and resulted in additional testing and analytical work late in the development program. The solution was to put an accumulator in the propellant feedline to reduce propellant oscillations.

The Space Shuttle Program took a proactive approach with a “Pogo Prevention Plan” drafted in the early 1970s. The plan called for comprehensive stability analysis and testing programs. Testing consisted of modal tests to verify the structural dynamic characteristics, hydroelastic tests of External Tank and propellant lines, and pulse testing of the Space Shuttle Main Engines. The plan baselined a Pogo suppression system—the first NASA launch vehicle to have such



a feature. The space agency selected and included an accumulator in the design of the main engines. This approach proved successful. Flight data demonstrated that the Space Shuttle was free of Pogo.



Test rig surrounds the Orbiter structural test article, Challenger, at the Lockheed Test Facility in Palmdale, California.

After the limit plus tests, the forward fuselage of the structural test article was subjected to a thermal environment gradient test. This testing entailed selective heating of the external skin regions with 25 zones. Gaseous nitrogen provided cooling. NASA used the data to assess the effects of thermal gradients and assist in the certification of thermal stresses by analysis techniques. Finally, the aft fuselage of the structural test article was subjected to internal/external pressures to provide strain and deflection data to verify the structural adequacy of the aft bulkhead and engine heat shield structures.

The structural test article subjected the Orbiter airframe to approximately 120% of limit load. To address ultimate load (140%) in critical areas, NASA conducted a series of supplemental tests on two major interfaces and 34 component specimens. The agency chose these specimens based on criticality of failure, uncertainty in analysis, and minimum fatigue margin. Designated specimens were subjected

to fatigue testing and analysis to verify the 100-mission life requirement. Finally, NASA tested all components to ultimate load and gathered data to compare predictions.

This unprecedented approach was challenged by NASA Headquarters and reviewed by an outside committee of experts from the “wide body” commercial aircraft industry. The experts concurred with the approach.

Acoustic Fatigue Integrity


Commercial and military aircraft commonly have a design life of 20,000 hours of flight composed of thousands of take offs and landings. As a result, the fatigue life is a design factor. The Orbiter, on the other hand, had a design life of 100 missions and a few hundred hours of flight in the atmosphere, but the acoustic environment during ascent was very high. Certification of acoustic fatigue life had to be accomplished.

The challenge was to certify this large, complex structure for a substantial number of combined acoustic, mechanical, and thermal conditions. No existing test facilities could accommodate a test article the size of the Orbiter or simulate all of the loads and environments.

The acoustic fatigue certification program was as innovative as that of the ultimate strength certification. The approach was to test a representative structure of various forms, materials, and types of construction in representative acoustic environments until the structure failed. This

Orbiter Acoustic Fatigue Test Articles

These acoustic fatigue test articles (shaded in blue) are representative of structure and environment.




Nozzle Flexible Bearing—Steering the Reusable Solid Rocket Motor

At Space Shuttle liftoff, initial steering was controlled in large part by the reusable solid rocket motors' movable nozzles. Large hydraulic actuators were attached to each nozzle. On command, these actuators mechanically vectored the nozzle, thereby redirecting the supersonic flow of hot gases from the motor.

A flexible bearing allowed the nozzle to be vectored. At about 2.5 m (8 ft) in diameter and 3,200 kg (7,000 pounds), this bearing was the largest flexible bearing in existence. The component had to vector up to 8 degrees while maintaining a pressure-tight seal against the combusive gases within the rocket, withstand high loads imparted at splashdown, and fit within the constraints of the solid rocket motor case segments. It also had to be reusable up to nine times.

The structure consisted of alternating layers of natural rubber (for flexibility) and steel shims (for strength and stiffness). The layers were spherically shaped,



During the first minutes of flight, a Thrust Vector Control System housed at the base of each solid rocket motor provided a majority of the steering capability for the shuttle. A flexible bearing enabled nozzle movement. Two hydraulic actuators generated the mechanical force needed to move the nozzle.

allowing the nozzle to pivot in any direction. Forces from the actuators induced a torque load on the bearing that strained the rubber layers in shear, with each layer rotating a proportional part of the total vector angle. This resulted in a change in nozzle angular direction relative to the rocket motor centerline.

The most significant manufacturing challenge was producing a vulcanization bond between the rubber and the shims.

Fabrication involved laying up the natural rubber by hand between the spherically shaped shims. Vulcanization was accomplished by applying pressure while controlling an elevated temperature gradient through the flexible bearing core. This process cured the rubber and vulcanized it to the shims in one step. The completed bearing underwent rigorous stretching and vectoring tests, including testing after each flight, as part of the refurbishment process.

established the level of damage that would be allowed for each type of structure. NASA selected 14 areas of the Orbiter to represent the various structural configurations.

The allowable damage was reduced analytically to account for the damage induced by the flight loads and temperature cycles for all regions of the vehicle.

Because of the high fatigue durability of the graphite-epoxy construction of the payload bay doors and Orbital Maneuvering System pods, these structures were not tested to failure. Instead, the strains measured during the acoustic tests were correlated with mathematical models and adequate fatigue life was demonstrated analytically. These test articles were subsequently used as flight hardware.

Summary

The unique approaches taken during the Space Shuttle Program in validating the structural integrity of the Orbiter airframe set a precedent in the NASA programs that followed. Even as more accurate analysis software and faster computers are developed, the need for anchoring predictions in the reality of testing remains a cornerstone in the safe flight of all space vehicles.

Pressure Vessel Experience

In the 1970s, NASA made an important decision—one based on previous experience and emerging technology—that would result in significant weight savings for shuttle. The agency implemented the Composite Overwrapped Pressure Vessels Program over the use of all-metal designs for storing high-pressure gases, 2,068 – 3,361 N/cm² (3,000 – 4,875 psi) oxygen, nitrogen, and helium. The agency used 22 such vessels in the Environmental Control and Life Support System, Reaction Control System, Main Propulsion System, and Orbital Maneuvering System. The basic new design consisted of a gas or liquid impermeable, thin-walled metal liner wrapped with a composite overwrap for primary pressure containment strength.

Safety—Always a Factor

The Space Shuttle Program built on the lessons learned from the Apollo Program. The pressure vessels were constructed of titanium and designed such that the burst pressure was only 1.5 times the operating pressure (safety factor). This safety factor was unprecedented at the time. To assure the safety of tanks with such a low margin of safety, NASA developed a robust qualification and acceptance program. The technical knowledge gained during the Apollo Program was leveraged by the shuttle, with the added introduction of a new type of pressure vessel to further reduce mass.

The Brunswick Corporation, Lake Forest, Illinois, developed, for the shuttle, a composite overwrapped

pressure vessel for high-pressure oxygen, nitrogen, and helium storage. The metallic liners were made of titanium (Inconel® for the oxygen systems) overwrapped with DuPont™ Kevlar® in an epoxy matrix. Switching from solid titanium tanks to composite overwrapped pressure vessels reduced the Space Shuttle tank mass by approximately 209 kg (460 pounds).


Since the shuttle was reusable and composite overwrapped pressure vessels were a new technology, the baseline factor of safety was 2.0. As development progressed, NASA introduced and instituted a formal fracture control plan based on lessons learned in the Apollo Program. As the composite overwrapped pressure vessels were fracture-critical items—e.g., their failure would lead to loss of vehicle and crew—fracture control required extensive lifetime testing of the vessels to quantify all failure modes. The failure mechanisms of the composite were just beginning to be understood. Kevlar® is very durable, so minor damage to the overwrap was not critical. NASA, however, discovered

that the composite could fail when under a sustained stress, less than its ultimate capability, and could fail without indication. This failure mode of the composite was called “stress rupture” and could lead to a catastrophic burst of the pressure vessel since the metallic liner could not carry the pressure stress alone.

In the late 1970s, engineers observed unexpectedly poor stress rupture performance in the testing of Kevlar® strands at the Lawrence Livermore National Laboratory in Livermore, California. As a result, NASA contracted with that laboratory to study the failure modes of the Kevlar® fiber for application in the shuttle tanks. Technicians conducted hundreds of tests on individual Kevlar® fibers, fiber/epoxy strands, and subscale vessels.

The development program to characterize all the failure modes of the composite overwrapped pressure vessels set the standard for all spaceflight programs. Therefore, as tank development proceeded, NASA used the fracture control test program to

Composite Overwrapped Pressure Vessels



NASA Puts Vessels to the “Stress Test”

In 1978, NASA developed and implemented a “fleet leader” test program to provide Orbiter subscale vessel stress rupture data for comparison to existing strand and subscale vessel data. Vessels in the test program were subscale in size and used aluminum liners instead of titanium, yet they were built by the same company manufacturing the Orbiter composite overwrapped pressure vessels using the same materials, equipment, and processes/procedures. These vessels were put to test at Johnson Space Center in Houston, Texas.

The test program consisted of two groups of vessels—15 vessels tested at ambient temperature conditions and an approximate stress level of 50% of ultimate strength; and 10 vessels tested at approximately 50% of average strength and an elevated temperature in an attempt to accelerate stress rupture failure. For the elevated temperature testing, 79°C (175°F) was



chosen as the test temperature for both groups. Engineers performed periodic depressurizations/repressurizations to simulate Orbiter usage and any potential effects.

The ambient temperature vessels were pressurized for nearly 25 years without failure before NASA stopped testing. The flight vessels only accumulated a week or two worth of pressure per

mission, so the ground tests led the fleet by a significant margin.

For the accelerated 79°C (175°F) temperature testing, the first failure occurred after approximately 12 years and the second at 15 years of pressure. These stress rupture failures indicated that the original stress rupture life predictions for composite overwrapped pressure vessels were conservative.

justify a safe reduction in the factor of safety on burst from 2.0 to 1.5, resulting in an additional 546 kg (1,203 pounds) of mass saved from the Orbiter.

Even with all of the development testing, two non-stress rupture composite overwrapped pressure vessels failures occurred on shuttle. The complexity of the welding process on certain materials contributed to these failures. To build a spherical

pressure vessel, two titanium hemispheres had to be welded together to form the liner. Welding titanium is difficult and unintentional voids are sometimes created. Voids in the welds of two Main Propulsion System vessels had been missed during the acceptance inspection. In May 1991, a Main Propulsion System helium pressurization vessel started leaking on the Atlantis prior to the launch of

Space Transportation System (STS)-43. NASA removed these vessels from the Orbiter.

The subsequent failure investigation found that, during manufacture, 89 pores formed in the weld whereas the typical number for other Orbiter vessels was 15. Radiographic inspection of the welds showed that the pores had initiated fatigue cracks that eventually broke through the liner, thereby causing



the leak. While this inspection was ongoing, the other Main Propulsion System vessel on Atlantis started leaking helium—once again due to weld porosity. NASA reviewed all other vessels in service, but none had weld porosity levels comparable to the two vessels that had leaked.

Space Shuttle Experiences Influence Future Endeavors

NASA's Orbiter Project pushed the technology envelope for pressure vessel design. Lessons learned from development, qualification, and in-service failures prompted the International Space Station (ISS) and future space and science missions to develop more robust requirements and verification programs. The ISS Program instituted structure controls based on the shuttle investigation of pressure vessels. No other leaks in pressure vessel tanks occurred through 2010—STS-132. For instance, the factor of safety on burst pressure was 1.5; damage tolerance of the composite and metallic liner was clearly addressed through qualification testing and operational damage control plans; radiographic inspection of liner welds was mandatory with acceptable levels of porosity defined; and material controls were in place to mitigate failure from corrosion, propellant spills, and stress rupture. These industry standard design requirements for composite overwrapped pressure vessels are directly attributable to the shuttle experience as well as its positive influence on future spaceflight.

Fracture Control Technology Innovations— From the Space Shuttle Program to Worldwide Use

A fundamental assumption in structural engineering is that all components have small flaws or crack-like defects that are introduced during manufacturing or service. Growth of such cracks during service can lead to reduced service life and even catastrophic structural failure. Fracture control methodology and fracture mechanics tools are important means for preventing or mitigating the adverse effects of such cracks. This is important for industries where structural integrity is of paramount importance.

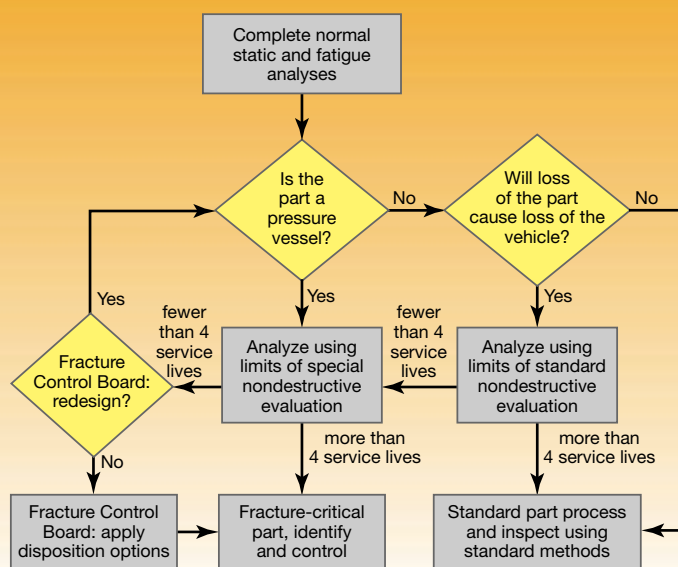
Prior to the Space Shuttle, NASA did not develop or implement many fracture mechanics and fracture control applications during the design and build phases of space vehicles. The prevailing design philosophy at the time was that safety factors on static strength provided a margin against fracture and that simple proof tests of tanks (pressure vessels) were sufficient to demonstrate the margin of safety. In practice, however, the Apollo Program experienced a number of premature test failures of pressure vessels that resulted in NASA implementing a version of fracture control referred to as “proof test logic.” It was not until the early 1960s that proof tests were sufficiently understood from a fracture mechanics point of view—that proof tests could actually be used, in some cases, to ensure the absence of initial flaws of a size that could cause failure within a pressure vessel's operating conditions.

The application of proof test logic required the determination of environmental crack growth thresholds for all environments to which the pressure vessels were exposed while pressurized as well as development of fracture toughness values and cyclic crack growth rates for materials used in the pressure vessels. The thresholds resulted in pressurization restrictions and environmental control of all Apollo pressure vessels. In effect, proof test logic formed the first implementation of a rigorous fracture control program in NASA.

Fracture Control Comes of Age

The legacy of the Apollo pressure vessel failure experience was that NASA, through the Space Shuttle Program, became an industry leader in the development and application of fracture mechanics technology and fracture control methodology. Although proof test logic worked successfully for the Apollo pressure vessels, the Space Shuttle Program brought with it a wide variety of safety-critical, structurally complex components (not just pressure vessels), materials with a wide range of fracture properties, and an aircraft-like fatigue environment—all conditions for which proof test logic methodology could not be used for flaw screening purposes. The shuttle's reusable structure demanded a more comprehensive fracture control methodology. In 1973, the Orbiter Project released its fracture control plan that set the requirements for and helped guide the Orbiter hardware through the design and build phases of the project.

How NASA Determined What Parts Required Attention



Early Shuttle Fracture Control

Fracture control, as practiced early in the Space Shuttle Program, was a three-step process: select the candidate fracture critical components, perform fracture mechanics analyses of the candidates, and disposition the components that had insufficient life.

Design and stress engineers selected the candidate fracture critical components. The selection was based on whether failure of the component from crack propagation could lead to a loss of life or vehicle. Certain components, such as pressure vessels, were automatically considered fracture critical. Performing a fracture mechanics analysis of the candidates started with an assumed initial crack located in the most unfavorable location in the component. The size of the assumed crack was typically based on the nondestructive inspection that was performed on the component. The fracture mechanics analysis

required knowledge of the applied stress, load spectrum, environment, assumed initial crack size, materials fracture toughness, and materials fatigue and environmental crack growth properties. Fracture analysis was required to show a service life of four times the shuttle's 100-mission design life.

There were a number of options for dispositioning components that had insufficient life. These options included the following:

- Redesigning the component when weight and cost permitted
- Conducting nondestructive inspection with a more sensitive technique where special nondestructive evaluation procedures allowed a smaller assumed crack size
- Limiting the life of the component
- Considering multiple element load paths
- Demonstrating life by fracture mechanics testing of the component

- Refining the loading based on actual measurements from the full-scale structural test articles

In addition to being a fundamental part of the structural design process, fracture mechanics became a useful tool in failure analysis throughout the Space Shuttle Program.

Fracture Control Evolves with Payloads

The shuttle payload community further refined the Orbiter fracture control requirements to ensure that a structural failure in a payload would not compromise the Space Shuttle or its Orbiter. NASA classified payloads by the nature of their safety criticality. Typically, a standard fracture criticality classification process started by removing all exempt parts that were nonstructural items—i.e., items not susceptible to crack propagation such as insulation blankets or certain common small parts with well-developed quality-control programs and use history.

All remaining parts were then assessed as to whether they could be classified as non-fracture critical. This category included the following classifications:

- Low-released mass—parts with a mass low enough that, if released during a launch or landing, would cause no damage to other components
- Contained—a failed part confined in a container or otherwise restrained from free release
- Fail-safe—structurally redundant designs where remaining components could adequately and safely sustain the loading that the failed member would have carried or failure would not result in a catastrophic event
- Low risk—parts with large structural margins or other conditions making crack propagation extremely unlikely

- Nonhazardous leak-before-burst—pressure vessels that did not contain a hazardous fluid where loss of fluid would not cause a catastrophic hazard such as loss of vehicle and crew, and where the critical crack size was much greater than the vessel wall thickness

NASA processed non-fracture critical components under conventional aerospace industry verification and quality assurance procedures.

All parts that could not be classified as exempt or non-fracture critical were classified as fracture critical. Fracture critical components had to have their damage tolerance demonstrated by testing or by analysis. To assure conservative results, such tests or analyses assumed that a flaw was located in the most unfavorable location and was subjected to the most unfavorable loads. The size of the assumed flaw was based on the nondestructive inspections that were used to inspect the hardware. The tests or analyses had to demonstrate that such an assumed crack would not propagate to failure within four service lifetimes.

Fracture Control Software Development

Few analytical tools were available for fracture mechanics analysis at the start of the Space Shuttle Program. The number of available analytical solutions was limited to a few idealized crack and loading configurations, and information on material dependency was scarce. Certainly, computing power and availability provided no comparison to what eventually became available to engineers. Improved tools to effect the expanded application of fracture mechanics and fracture control were deemed necessary for safe operation of the shuttle.

With Space Shuttle Program support, Johnson Space Center (JSC) initiated a concerted effort in the mid 1970s to create a comprehensive database of materials fracture properties. This involved testing virtually all metallic materials in use in the program for their fracture toughness, environmental crack growth thresholds, and fatigue crack growth rate properties. NASA manufactured and tested specimens in the environments that Space Shuttle components experienced—cryogenic, room, and elevated temperatures as well as in vacuum, low- and high-humidity air, and selected gaseous or fluid environments. Simultaneously, a parallel program created a comprehensive library of analytical solutions. This involved compiling the small number of known solutions from various sources as well as the arduous task of deriving new ones applicable to shuttle configurations.


Fatigue Crack Computer Program

By the early 1980s, JSC engineers developed a computer program—NASA/FLAGRO—to provide fracture data and fracture analysis for crewed and uncrewed spacecraft components. NASA/FLAGRO was the first known program to contain comprehensive libraries of crack case solutions, material fracture properties, and crack propagation models. It provided the means for efficient and accurate analysis of fracture problems.

NASGRO® Becomes a Worldwide Standard in Fracture Analysis

Although NASA/FLAGRO was essentially a shuttle project, NASA eventually formed an agencywide fracture control methodology panel to standardize fracture methods and requirements across the agency and to guide the development of

Crack Models and Material Properties Required for Fracture Analyses




Crack in a payload mounting plate.



Fracture mechanics pretest and posttest specimens for characterizing material behavior.

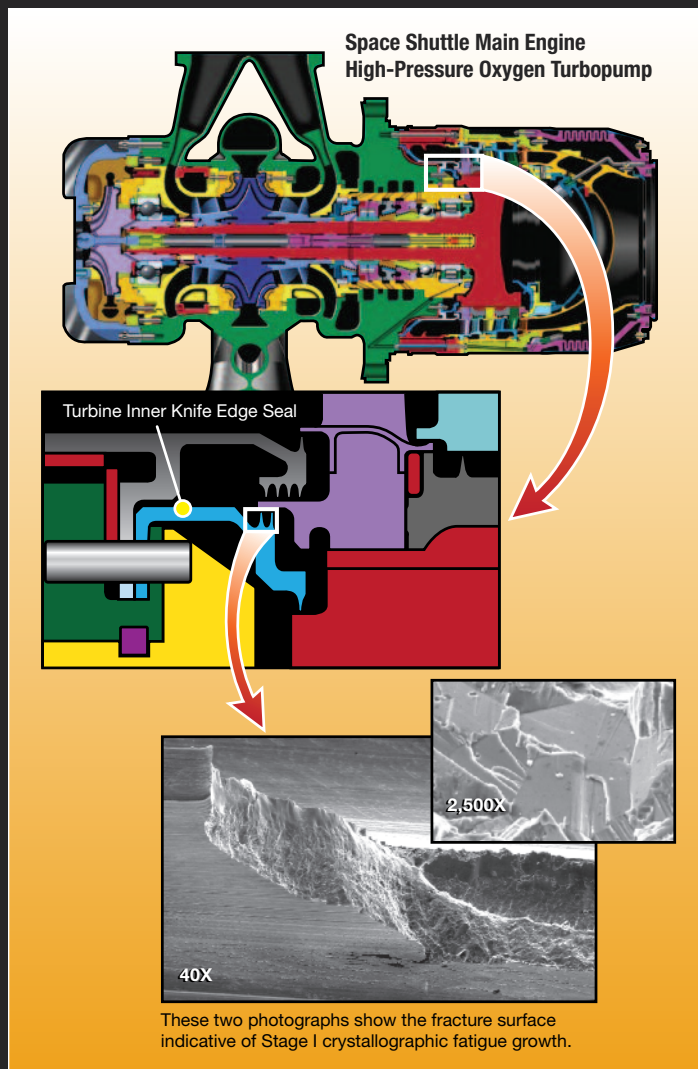
Typical NASGRO® analytical model of cracked structure for prediction of fatigue and fracture behavior, in which the crack driving force (K) is a function of the applied stress (σ) and the crack depth (a).



Space Shuttle Main Engine Fracture Control

The early Space Shuttle Main Engine (SSME) criteria for selecting fracture critical parts included Inconel® 718 parts that were exposed to gaseous hydrogen. These specific parts were selected because of their potential for hydrogen embrittlement and increased crack growth caused by such exposure. Other parts such as turbine disks and blades were included for their potential to produce shrapnel. Titanium parts were identified as fracture critical because of susceptibility to stress corrosion cracking. Using these early criteria, approximately 59 SSME parts involving some 290 welds were identified as being fracture critical.

By the time the alternate turbopumps were introduced into the shuttle fleet in the mid 1990s, fracture control processes had been well defined. Parts were identified as fracture critical if their failure due to cracking would result in a catastrophic event. The fracture critical parts were inspected for preexisting cracks, a fracture mechanics assessment was performed, and materials traceability, and part-specific life limits were imposed as necessary. This combination of inspection, analysis, and life limits ensured SSME fracture critical parts were flown with confidence.




NASA/FLAGRO, renamed NASGRO®, for partnership with industry. While other commercial computer programs existed by the end of the Space Shuttle Program, none had approached NASGRO® in its breadth of technical capabilities, the size of its fracture solution library, and the size of its materials database. In addition to gaining several prestigious engineering awards, NASGRO® is in use by organizations and companies around the world.

Summary

Fracture mechanics is a technical discipline first used in the Apollo Program, yet it really came of age in the Space Shuttle Program. Although there is still much to be learned, NASA made great strides in the intervening 4 decades of the shuttle era in understanding the physics of fracture and the methodology of fracture control. It was this agency's need to analyze shuttle and payload fracture critical structural hardware that led to the

development of fracture mechanics as a tool in fracture control and ultimately to the development of NASGRO®—the internationally recognized fracture mechanics analysis software tool. The shuttle was not only a principal benefactor of the development of fracture control, it was also the principal sponsor of its development.



Robotics and Automation

Introduction

Gail Chapline
Steven Sullivan

Shuttle Robotic Arm

Henry Kaupp
Elizabeth Bains
Rose Flores
Glenn Jorgensen
Y.M. Kuo
Harold White

Automation: The Space Shuttle Launch Processing System

Timothy McKelvey

Integrated Network Control System

Wayne McClellan
Robert Brown

Orbiter Window Inspection

Bradley Burns

Robotics System Sprayed Thermal Protection on Solid Rocket Booster

Terry Huss
Jack Scarpa

Although shuttle astronauts made their work in space look like an everyday event, it was in fact a hazardous operation. Using robotics or human-assisted robotics and automation eliminated the risk to the crew while still performing the tasks needed to meet the mission objectives. The Shuttle Robotic Arm, commonly referred to as “the arm,” was designed for functions that were better performed by a robotic system in space.

Automation also played an important role in ground processing, inspection and checkout, cost reduction, and hazardous operations. For each launch, an enormous amount of data from verification testing, monitoring, and command procedures were compiled and processed, often simultaneously. These procedures could not be done manually, so ground automation systems were used to achieve accurate and precise results. Automated real-time communication systems between the pad and the vehicle also played a critical role during launch attempts. In addition, to protect employees, automated systems were used to load hazardous commodities, such as fuel, during tanking procedures. Throughout the Space Shuttle Program, NASA led the development and use of the most impressive innovations in robotics and automation.



Shuttle Robotic Arm— Now That You Have the “TRUCK,” How Do You Make the Delivery?

Early in the development of the Space Shuttle, it became clear that NASA needed a method of deploying and retrieving cargo from the shuttle payload bay. Preliminary studies indicated the need for some type of robotic arm to provide both capabilities. This prompted the inclusion of a Shuttle Robotic Arm that could handle payloads of up to 29,478 kg (65,000 pounds).

In December 1969, Dr. Thomas Paine, then administrator of NASA, visited Canada and extended an offer for Canadian participation with a focus on the Space Shuttle. This was a result of interest by NASA and the US government in foreign participation in post-Apollo human space programs. In 1972, the Canadian government indicated interest in developing the Shuttle Robotic Arm. In 1975, Canada entered into an agreement with the US government in which Canada would build the robotic arm that would be operated by NASA.

The Shuttle Robotic Arm was a three-joint, six-degrees-of-freedom, two-segment manipulator arm to be operated only in the microgravity

environment. From a technical perspective, it combined teleoperator technology and composite material technology to produce a lightweight system useable for space applications. In fact, the arm could not support its own weight on Earth. The need for a means of grapppling the payload for deployment and retrieval became apparent. This led to an end effector—a unique electromechanical device made to capture payloads.

Unique development and challenges of hardware, software, and extensive modeling and analysis went into the Shuttle Robotic Arm's use as a tool for delivery and return of payloads to and from orbit. Its role continued in the deployment and repair of the Hubble



Backdropped by the blackness of space and Earth's horizon, Atlantis' Orbiter Docking System (foreground) and the Canadarm—the Shuttle Robotic Arm developed by Canada—in the payload bay are featured in this image photographed by an STS-122 (2008) crew member during Flight Day 2 activities.

Space Telescope, its use in the building of the space station and, finally, in Return to Flight as an inspection and repair tool for the Orbiter Thermal Protection System.

Evolution of the Shuttle Robotic Arm

The initial job of the Shuttle Robotic Arm was to deploy and retrieve payloads to and from space. To accomplish this mission, the system that was developed consisted of an anthropomorphic manipulator arm located in the shuttle cargo bay, cabin equipment to provide an interface to the main shuttle computer, and a human interface to allow an astronaut to control arm operations remotely.


The manipulator arm consisted of three joints, two arm booms, an end effector, a Thermal Protection System, and a closed-circuit television system. Arm joints included a shoulder joint with two degrees of freedom (yaw and pitch), an elbow joint with one degree of freedom (pitch), and a wrist joint with three degrees of freedom (pitch, yaw, and roll). Each joint degree of freedom consisted of a motor module driving a gear box to effect joint movement and appropriate local processing to interpret drive commands originating from the cabin electronics.

The cabin electronics consisted of a displays and controls subsystem that provided the human-machine interface to allow a crew member to command the arm and display appropriate information, including arm position and velocity, end effector status, temperature, and caution and warning information. Additionally, in the displays and controls subsystem, two hand controllers allowed man-in-the-loop control of the end point of the

arm. The main robotic arm processor—also part of the cabin electronics—handled all data transfer among the arm, the displays and controls panel, and the main shuttle computer. The main shuttle computer processed commands from the operator via the displays and controls panel; received arm data to determine real-time position, orientation, and velocity; and then generated rate and

current limit commands that were sent to the arm-based electronics.


The arm was thermally protected with specially designed blankets to reduce the susceptibility of the hardware to thermal extremes experienced during spaceflight and had an active thermostatically controlled and redundant heater system.





Components of the Shuttle Robotic Arm

Crew Compartment



With a total length of 15.24 m (50 ft), the Shuttle Robotic Arm consisted of two lightweight high-strength tubes, each 0.381 m (1.25 ft) in diameter and 6.71 m (22 ft) in length, with an elbow joint between them. From a shoulder joint at the base of the arm providing yaw and pitch movement, the upper boom extended outward to the elbow joint providing pitch movement from which the lower arm boom stretched to a wrist joint providing pitch, yaw, and roll movement. The end effector was used to grapple the payload.

The closed-circuit television system consisted of a color camera on a pan/tilt unit near the elbow joint and a second camera in a fixed location on the wrist joint, which was primarily used to view a grapple fixture target when the arm was capturing a payload.


Self checks existed throughout all the Shuttle Robotic Arm electronics to assess arm performance and apply appropriate commands to stop the arm, should a failure occur. Caution and warning displays provided the operator with insight into the cause of the failure and remaining capability to facilitate the development of a workaround plan.

The interfacing end of the Shuttle Robotic Arm was equipped with a fairly complicated electromechanical construction referred to as the end effector. This device, the analog to a human hand, was used to grab, or grapple, a payload by means of a tailored interface known as a grapple fixture.

The end effector was equipped with a camera and light used to view the grapple fixture target on the payload being captured. The robotic arm provided video to the crew at the aft flight deck, and the camera view helped the crew properly position the end


Close-up View of End Effector and Grapple Fixture





© MacDonald, Detwiler and Associates Ltd. All rights reserved.

End Effector Capture/Rigidize Sequence: The left frame illustrates the snares in the open configuration, and the second frame shows the snares closed around the grapple shaft and under the grapple cam at the tip of the grapple shaft. The next frame illustrates the snares pulling the grapple shaft inside the end effector so the three lobes are nested into the mating slots in the end effector, and the final frame shows the snare cables being pulled taut to ensure a snug interface that could transfer all of the loads.



© MacDonald, Detwiler and Associates Ltd. All rights reserved.

Flat floor testing of the Shuttle Robotic Arm.



Challenger's (STS-8 [1983]) payload flight test article is lifted from the payload bay and held over clouds and water on Earth.

effector relative to the grapple fixture prior to capturing a payload. When satisfied with the relative position of the end effector to the payload grapple fixture using the grapple fixture target, the crew executed a command to capture and secure the payload.

Since the Shuttle Robotic Arm could not lift its own weight on Earth, all proposed operations had to be tested with simulations. In fact, terrestrial certification was a significant engineering challenge. Developing the complex equations describing the six-degrees-of-freedom arm was

one technical challenge, but solving equations combining 0.2268-kg (0.5-pound) motor shafts and 29,478-kg (65,000-pound) payloads also challenged computers at the time. Canada—the provider of the Shuttle Robotic Arm—and the United States both developed simulation models. The simulation responses were tested against each other as well as data from component tests (e.g., motors, gearboxes) and flat floor tests. Final verification could be completed only on orbit. During four early shuttle flights, strain gauges were added to the Shuttle Robotic Arm to measure loads during

test operations that started with an unloaded arm and then tested the arm handling progressively heavier payloads up to one emulating the inertia of a 7,256-kg (16,000-pound) payload—the payload flight test article. These data were used to verify the Shuttle Robotic Arm models.

Future on-orbit operations were tested preflight in ground-based simulations both with and without an operator controlling the Shuttle Robotic Arm. Simulations with an operator in the loop used mock-ups of the shuttle cockpit and required calculation of arm




response between the time the operator commanded arm motion with hand controllers or computer display entries and the time the arm would respond to commands on orbit. This was a significant challenge to then-current computers and required careful simplification of the arm dynamics equations. During the late 1970s and early 1980s, this necessitated banks of computers to process dynamic equations and specialized computers to generate the scenes. The first electronic scene generator was developed for simulations of shuttle operations, and

payload handling simulations drove improvements to this technology until it became attractive to other industries. Simulations that did not require an operator in the loop were performed with higher complexity equations. This allowed computation of loads within the Shuttle Robotic Arm and detailed evaluation of performance of components such as motors.

Since the Shuttle Robotic Arm's job was to deploy and retrieve payloads to and from space, NASA determined two cameras on the elbow and wrist would be invaluable for mission support

viewing since the arm could be maneuvered to many places the fixed payload bay cameras could not capture. As missions and additional hardware developed, unique uses of the arm emerged. These included "cherry picking" in space using a mobile foot restraint that allowed a member of the crew to have a movable platform from which tasks could be accomplished; "ice busting" to remove a large icicle that formed on the shuttle's waste nozzle; and "fly swatting" to engage a switch lever on a satellite that had been incorrectly positioned.




Astronauts Joseph Acaba and Akihiko Hoshide in the functional shuttle aft cockpit in the Systems Engineering Simulator showing views seen out of the windows. The Systems Engineering Simulator is located at NASA Johnson Space Center, Houston, Texas.



Cherry picking—On STS-41B (1984), Astronaut Bruce McCandless tests a mobile foot restraint attached to the Shuttle Robotic Arm. This device, which allowed a crew member to have a movable platform in space from which tasks could be accomplished, was used by shuttle crews throughout the program.



Fly swatting—On STS-51D (1985), the spacecraft sequencer on the Leasat-3 satellite failed to initiate antenna deployment, spin-up, and ignition of the perigee kick motor. The mission was extended 2 days to make the proper adjustments. Astronauts David Griggs and Jeffrey Hoffman performed a spacewalk to attach “fly swatter” devices to the robotic arm. Rhea Seddon engaged the satellite’s lever using the arm and the attached “fly swatter” devices.



Ice busting—On STS-41D (1984), a large icicle formed on the shuttle's waste nozzle. NASA decided that the icicle needed to be removed prior to re-entry into Earth's atmosphere. The Shuttle Robotic Arm, controlled by Commander Henry Hartsfield, removed the icicle.

The Hubble Missions

The Hubble Space Telescope, deployed on Space Transportation System (STS)-31 (1990), gave the world a new perspective on our understanding of the cosmos. An initial problem with the telescope led to the first servicing mission and the desire to keep studying the cosmos. The replacement and enhancement of the instrumentation led to a number of other servicing missions: STS-61 (1993), STS-82 (1997), STS-103 (1999), STS-109 (2002), and STS-125 (2009). From a Shuttle Robotic Arm perspective, the Hubble servicing missions showcased the system's ability to capture, berth, and release a relatively large payload as well as support numerous spacewalks to complete repair and refurbishment activities.


In the case of Hubble, the crew captured and mated the telescope to a berthing mechanism mounted in the payload bay to facilitate the repair and refurbishment activities. In this

scenario, a keel target mounted to the bottom of Hubble was viewed with a keel camera and the crew used the Shuttle Robotic Arm to position the Hubble properly relative to its berthing interface to capture and latch it.


The Era of Space Station Construction

With STS-88 (1998)—the attachment of the Russian Zarya module to the space station node—the attention of the shuttle and, therefore, the Shuttle Robotic Arm was directed to the construction of the space station. Early space station flights can be divided broadly into two categories: logistics flights and construction flights. With the advent of the three Italian-built Multi-Purpose Logistic Modules, the Shuttle Robotic Arm was needed to berth the modules to the station. The construction flights meant attaching a new piece of hardware to the existing station. Berthings were used to install new elements: the nodes; the modules, such as the US Laboratory Module and the Space Station Airlock; the truss segments, many of which contained solar panels for power to the station; and the Space Station Robotic Arm. These activities required some modifications to the Shuttle Robotic Arm as well as the addition of systems to enhance alignment and berthing operations.

During preliminary planning, studies evaluated the adequacy of the Shuttle Robotic Arm to handle the anticipated payload operations envisioned for the space station construction. These studies determined that arm controllability would not be satisfactory for the massive payloads the arm would need to manipulate.



A robotic vision system known as the Space Vision System was used for the first space station assembly flight (STS-88 [1988]) that attached Node 1 to the Russian module Zarya. This Space Vision System used a robotic vision algorithm to interpret relative positions of target arrays on each module to calculate the relative position between the two berthing interfaces. The crew used these data to enhance placement to ensure a proper berthing. The two panes above show the camera views from the shuttle payload bay that the robotic vision system analyzed to provide a relative pose to the crew.



Centerline Berthing Camera System: A Centerline Berthing Camera System was later adopted to facilitate ease of use and to enhance the ability of the crew to determine relative placement between payload elements. The left pane shows the centerline berthing camera mounted in a hatch window with its light-emitting diodes illuminated. The right pane shows the display the crew used to determine relative placement of the payload to the berthing interface. The outer ring of light-emitting diode reflections come from the window pane that the camera was mounted against. However, these reflections never moved and were ignored. The small ring at the center of the crosshairs is the reflection of the Centerline Berthing Camera System light-emitting diodes in the approaching payload window being maneuvered by the Shuttle Robotic Arm system. This was used to determine the angular misalignment (pitch and yaw) of the payload. The red chevrons to the left and right were used to determine vertical misalignment and roll while the top red chevron was used to determine horizontal misalignment. The green chevrons in the overlay were used to determine the range of the payload. This system was first used during STS-98 (2001) to berth the US Laboratory Module (Destiny) to Node 1.

Redesigning the arm-based electronics in each joint provided the necessary controllability. The addition of increased self checks also assured better control of hardware failures that could cause hazardous on-orbit conditions.

During the process of assembling the space station, enhanced berthing cue systems were necessary to mate complicated interfaces that would need to transmit loads and maintain a pressurized interior. The complexity and close tolerance of mating parts led to the development of several berthing

cue systems, such as the Space Vision System and the Centerline Berthing Camera System, to enhance the crew's ability to determine relative position between mating modules.


Return to Flight After Columbia Accident

During the launch of STS-107 (2003), a piece of debris hit the shuttle, causing a rupture in the Thermal Protection System that is necessary for re-entry into Earth's atmosphere, thereby leading to the Columbia accident. The ramifications of this breach in the shuttle's Thermal Protection System changed the role of the robotic arm substantially for all post-Columbia-accident missions. Development of the robotically compatible 15.24-m (50-ft) Orbiter Boom Sensor System provided a shuttle inspection and repair capability that addressed the Thermal Protection System inspection requirement for post-Columbia Return to Flight missions. Modification of the robotic arm wiring provided power and data capabilities to support inspection cameras and lasers at the tip of the inspection boom.

Two shuttle repair capabilities were provided in support of the Return to Flight effort. The first repair scenario required the Shuttle Robotic Arm, grappled to the space station, to position the shuttle and the space station in a configuration that would enable a crew member on the Space Station Robotic Arm to perform a repair. This was referred to as the Orbiter repair maneuver. The second repair scenario involved the Shuttle Robotic Arm holding the boom with the astronaut at the tip.



Orbiter Boom Sensor System



© MacDonald, Dettwiler and Associates Ltd. All rights reserved.

The operational scenario was that, post ascent and pre re-entry into Earth's atmosphere, the robotic arm would reach over to the starboard side and grapple the Orbiter Boom Sensor System at the forward grapple fixture and unberth it. The robotic arm and boom would then be used to pose the inspection sensors at predetermined locations for a complete inspection of all critical Thermal Protection System surfaces. This task was broken up into phases: inspect the starboard side, the nose, the crew cabin, and the port side. When the scan was complete, the robotic arm would berth the Orbiter Boom Sensor System back on the starboard sill of the shuttle and continue with mission objectives.


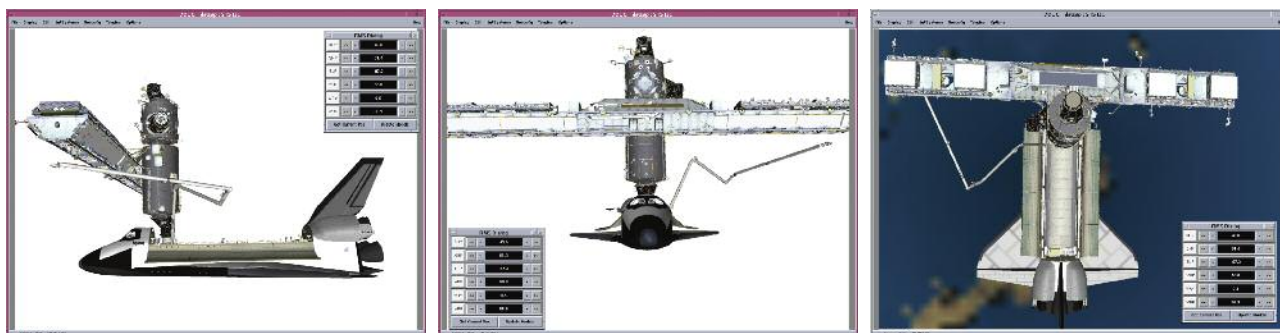


Image from STS-114 (2005) of the Orbiter Boom Sensor System scanning the Orbiter.

All post-Columbia-accident missions employed the Shuttle Robotic Arm and Orbiter Boom Sensor System combination to survey the shuttle for damage. The robotic arm and boom were used to inspect all critical Thermal Protection System surfaces. After the imagery data were processed, focused inspections occasionally followed to obtain additional images of areas deemed questionable from the inspection. A detailed test objective on STS-121 (2006) demonstrated the feasibility of having a crew member




Graphic simulation of Shuttle Robotic Arm/Orbiter Boom Sensor System-based repair scenario for port wing tip, starboard wing, and Orbiter aft locations.



Graphic simulation of the configuration of the Shuttle Robotic Arm/Orbiter Boom Sensor System for STS-121 (2006) flight test.

on the end of the combined system performing actions similar to those necessary for Thermal Protection System repair. Test results showed that the integrated system could be used as a repair platform and the system was controllable with the correct control parameters, good crew training, and proper extravehicular activity procedures development.

In support of shuttle repair capability and rescue of the crew, simulation tools were updated to facilitate the handling of both the space station and another shuttle as “payloads.” The space station as a payload was discussed earlier as a Return to Flight capability, known as the Orbiter repair maneuver. The shuttle as a payload came about due to the potential for a



In addition to performing inspections, the Orbiter Boom Sensor System’s role was expanded to include the ability to hold a crew in position for a repair to the Thermal Protection System. Considering that this was a 30.48-m (100-ft) robotic system, there was concern over the dynamic behavior of this integrated system. The agency decided to perform a test to evaluate the stability and strength of the system during STS-121 (2006).



Hubble rescue mission. Given that the space station would not be available for crew rescue for the final Hubble servicing mission, another shuttle would be “ready to go” on another launch pad in the event the first shuttle became disabled. For the crew from the disabled shuttle to get to the rescue shuttle, the Shuttle Robotic Arm would act as an emergency pole between the two vehicles, thus making the payload for the Shuttle Robotic Arm another shuttle. Neither of these repair/rescue capabilities—Orbiter repair maneuver or Hubble rescue—ever had to be used.

Summary

The evolution of the Shuttle Robotic Arm represents one of the great legacies of the shuttle, and it provided the impetus and foundation for the Space Station Robotic Arm. From the early days of payload deployment and retrieval, to the development of berthing aids and techniques, to the ability to inspect the shuttle for damage and perform any necessary repairs, the journey has been remarkable and will serve as a blueprint for space robotics in the future.

Automation: The Space Shuttle Launch Processing System

The Launch Processing System supported the Space Shuttle Program for over 30 years evolving and adapting to changing requirements and technology and overcoming obsolescence challenges.

Designed and developed in the early 1970s, the Launch Processing System began operations in September 1977 with a focused emphasis on safety, operational resiliency, modularity, and flexibility. Over the years, the system expanded to include several firing rooms and smaller, specialized satellite sets to meet the processing needs of multiple Space Shuttles—from landing to launch.

Architecture and Innovations

The architecture of the system and innovations included in the original design were major reasons for the Launch Processing System’s outstanding success. The system design required that numerous computers had the capability to share real-time measurement and status data with each other about the shuttle, ground support equipment, and the health and status of the Launch Processing System itself. There were no commercially available products to support the large-scale distributed computer network required for the system. The solution to this problem was to network the Space Shuttle firing room computers using a centralized hub of memory called a common data buffer—designed by NASA at Kennedy Space Center (KSC) specifically for computer-to-


computer communication. The buffer was a high-speed memory device that provided shared memory used by all command and control computers supporting a test. Each computer using the buffer was assigned a unique area of memory where only that computer could write data; however, every computer on the buffer could read those data. The buffer could support as many as 64 computers simultaneously and was designed with multiple layers of internal redundancy, including error-correcting software. The common data buffer’s capability to provide fast and reliable intercomputer communication made it the foundation of the command and control capability of the firing room.

The System Console

Other outstanding features of the Launch Processing System resided in the human-to-machine interface known as the console. System engineers used the console to control and monitor the particular system for which they were responsible. Each firing room contained 18 consoles—each connected to the common data buffer, and each supporting three separate command and control workstations. One of the key features of the console was its ability to execute up to six application software programs, simultaneously. Each console had six “concurrencies”—or areas in console memory—that could independently support an application program. This capability foreshadowed the personal computer with its ability to multitask using different windows. With six concurrencies available to execute as many as six application programs, the console operator could monitor



Launch Processing System



The Launch Processing System provides command and control of the flight vehicle elements and ground support equipment during operations at Kennedy Space Center.

thousands of pieces of information within his or her area of responsibility from a single location. Each console in the firing room was functionally identical, and each was capable of executing any set of application software programs. This meant any console could be assigned to support

any system, defined simply by what software was loaded. This flexibility allowed for several on-demand spare consoles for critical or hazardous tests such as launch countdown. The console also featured full color displays, programmable function keys, a programmable function

panel, full cursor control, and a print screen capability. Upgrades included a mouse, which was added to the console, and modernized cursor control and selection.

System Integrity

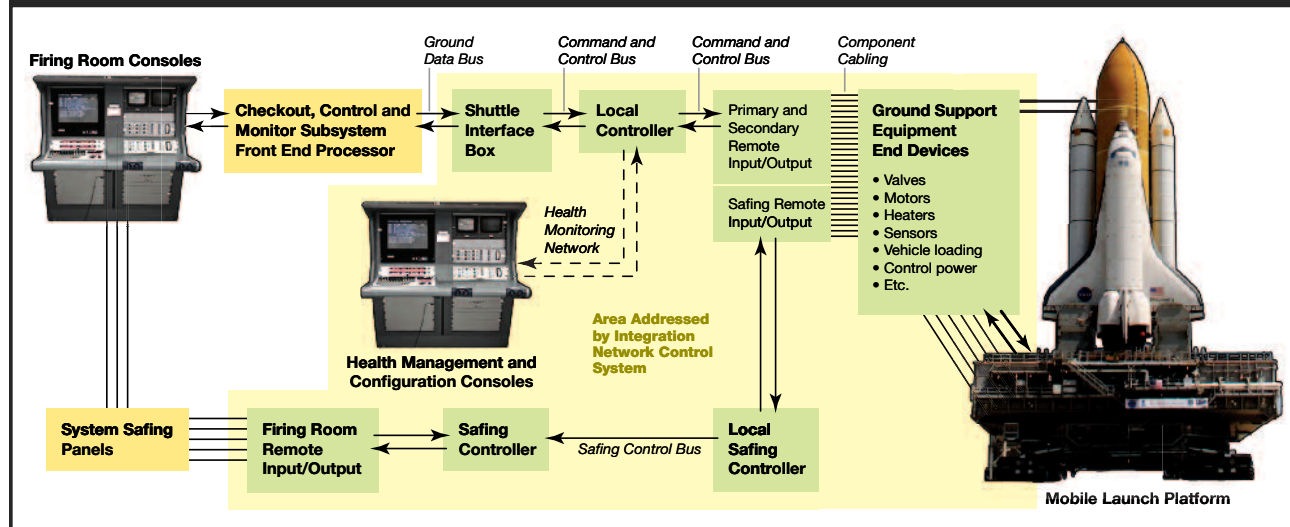
Fault tolerance, or the ability to both automatically and manually recover from a hardware or software failure, was designed and built into the Launch Processing System. An

equivalent analogy for distributed computer systems would be the clustering of servers for redundancy. Most critical computers within the system were operated in an active/standby configuration. A very high degree of system reliability

was achieved through automated redundancy of critical components.

A software program called System Integrity, which constantly monitored the health and status of all computers using the common data buffer,

Integrated Network Control System



The Integrated Network Control System was a reliable, automated network system that sent data and commands between the shuttle Launch Control Center and hardware end items. It bridged industry automation technologies with customized aerospace industry communication protocols and associated legacy end item equipment. The design met several challenges, including connectivity with 40,000 end items located within 28 separate ground systems, all dispersed to 10 facilities. It provided data reliability, integrity, and emergency safing systems to ensure safe, successful launch operations.

Ground control and instrumentation systems for the Space Shuttle Launch Processing System used custom digital-to-analog hardware and software connected to an analog wire-based distribution system. Loss of a data path during critical operations would compromise safety. To improve safety, data integrity, and

network connectivity, the Integrated Network Control System design used three independent networks.

The network topology used a quad-redundant, fiber-optic, fault-tolerant ring for long-distance distribution over the Launch Control Center, mobile launcher platforms, Orbiter processing facilities, and two launch pads. Shorter distances were accommodated with redundant media over coaxial cable for distribution over system and subsystem levels. This network reduced cable and wiring for ground processing over the Launch Complex 39 area by approximately 80% and cable interconnects by 75%. It also reduced maintenance and troubleshooting. This system was the first large-scale network control and health management system for the Space Shuttle Program and one of the largest, fully integrated control networks in the world.



governed the automatic recovery of failed critical computers in the firing room. In the event of a critical computer failure, System Integrity commanded a redundant switch, thereby shutting down the unhealthy computer and commanding the standby computer to take its place. Launch Processing System operators could then bring another standby computer on line from a pool of ready spares to reestablish the active/standby configuration.

Most critical portions of the Launch Processing System had redundancy and/or on-demand spare capabilities. Critical data communication buses between the Launch Control Center and the different areas where the shuttles were processed used both primary and backup buses. Critical ground support equipment measurements were provided with a level of redundancy, with a backup measurement residing on a fully independent circuit and processed by different firing room computers than the primary measurement. Electrical power to the firing room was supplied by dual uninterruptible power sources, enabling all critical systems to take advantage of two sources of uninterruptible power.

Critical software programs, such as those executed during launch countdown, were often part of the software load of two different consoles in the event of a console failure. The System Integrity program was executed simultaneously on two different firing room consoles. The fault tolerance designed into the Launch Processing System spanned from the individual measurement up through subsystem hardware and software, providing the Space Shuttle test team with outstanding operational resiliency in almost any failure scenario.

Orbiter Window Inspection

As the Orbiter moved through low-Earth orbit, micrometeors collided with it and produced hypervelocity impact craters that could produce weak points in its windows and cause the windows to fail during extreme conditions. Consequently, locating and evaluating these craters, as well as other damage, was critically important. Significant effort went into the development and use of ground window inspection techniques.



Bradley Burns, lead engineer in the development of the window inspection tool, monitors its progress as it scans an Orbiter window.

The window inspection tool could be directly attached to any of the six forward windows on any Orbiter. The tool consisted of a dual-camera system—a folded microscope and a direct stress imaging camera that was scanned over the entire area of the window. The stress imaging camera “saw” stress by launching polarized light at the window from an angle such that it bounced off the back of the window, then through the area being monitored, and finally into the camera where the polarization state was measured. Defects caused stress in the window. The stress changed the polarization of the light passing through it. The camera provided direct imaging of stress regions and, when coupled with the microscope, ensured the detection of significant defects.



The Portable Handheld Optical Window Inspection Device is vacuum attached to a window such that the small camera and optical sensor (black tube) were aimed at a defect.

The portable defect inspection device used an optical sensor. A three-dimensional topographic map of the defect could be obtained through scanning. Once a defect was found, the launch commit criteria was based on measuring the depth of that defect. If a window had a single defect deeper than a critical value, the window had to be replaced.



Robotics System Sprayed Thermal Protection on Solid Rocket Booster

Many Solid Rocket Booster (SRB) components were covered with a spray-on thermal protection material that shielded components from aerodynamic heating during ascent. The application process took place at the SRB Assembly and Refurbishment Facility at Kennedy Space Center. The process resulted in overspray and accounted for 27% of hazardous air emissions.

To address this drawback, NASA developed Marshall Convergent Coating-I, which consisted of improved mixing and robotic spray processes. The coating's ingredients were mixed (or *converged*) only during spraying. Hazardous waste was virtually eliminated after implementation of the system in the mid 1990s.

After each flight, the boosters were refurbished. This process began at NASA's Hangar AF Booster Recovery Facility at Cape Canaveral Air Force Station. There, a robotic high-pressure water jet, or "hydrolase," stripped the components of their Thermal Protection System materials.

NASA installed the hydrolase system in 1998. Each booster structure was numerically modeled. These models were used to program the robot to follow the contour of each component.

The Hangar AF wash facilities used a specially designed water filtration and circulation system to recycle and reuse the waste water.



An SRB aft skirt receives a robotically controlled layer of Marshall Convergent Coating-1 Thermal Protection System material.



A technician in a control booth monitors the robotic high-pressure hydrolase as it strips Thermal Protection System material from an SRB forward skirt.

Exception Monitoring

Another key concept designed into the Launch Processing System software was the capability to recognize and automatically react to out-of-bounds measurements. This capability was called exception monitoring, and it

monitored for specific measurements exceeding a predefined set of limits. When a Launch Processing System computer detected a measurement exception—for example, the pressure in a fuel tank exceeded its upper limit—the computer immediately notified the console responsible for that fuel tank.

A software program at the console promptly reacted to the exception and automatically sent a command or series of commands to resolve the problem. Similar software could also prevent inadvertent damage by verifying required parameters prior to command issuance, such as confirming that



pressures were appropriate prior to commanding a valve opening. Commands could also be manually sent by the console operator.

Survivability

Although the Launch Processing System's flexible architecture and distribution of hardware functionality allowed it to support the program consistently over 30 years, that support would not have been possible without a comprehensive and proactive sustaining engineering, maintenance, and upgrade approach. This is true for any large-scale computer system where an extended operational lifetime is desired.

The approach that kept the Launch Processing System operationally viable for over 3 decades was called the Survivability Program. Survivability was initiated to mitigate risk associated with the natural obsolescence of commercial off-the-shelf hardware products and the physical wear and tear on the electrical and mechanical subsystems within the Launch Processing System.

One of the main tenets of survivability was the desire to perform each upgrade with an absolutely minimal impact to system software. Hardware was upgraded to duplicate the existing hardware in form, fit, and function. The emphasis on minimizing software impacts was a distinct strength in survivability due to the resultant reduction of risk. Survivability projects were selected through careful analysis of maintenance failure data and constant surveillance of electronic manufacturers and suppliers by logistics to identify integrated circuits and other key components that were going to be unavailable in the near

future. Through this process, NASA purchased a "lifetime" buy of some electronic components and integrated circuits to ensure the Launch Processing System had ample spares for repair until the end of the program. It could also redesign a circuit board using available parts or replace an entire subsystem if a commercial off-the-shelf or in-house design solution offered the most benefit.

NASA eventually upgraded or replaced about 70% of the original Launch Processing System hardware under the survivability effort. The proactive application of the Survivability Program mitigated obsolescence and continued successful operational support.

Summary

These innovations and the distributed architecture of the Launch Processing System allowed upgrades to be performed over the years to ensure the system would survive through the life of the program. This success demonstrated that, with appropriate attention paid to architecture and system design and with proactive sustaining engineering and maintenance efforts, a large, modular, integrated system of computers could withstand the inevitable requirements change and obsolescence issues. It also demonstrated that it could successfully serve a program much longer than originally envisioned.

The Launch Processing System was vital to the success of KSC fulfilling its primary mission of flying out the Space Shuttle Program in a safe and reliable manner, thus contributing to the shuttle's overall legacy.



Systems Engineering for Life Cycle of Complex Systems

Introduction

Gail Chapline

Steven Sullivan

Systems Engineering During Development of the Shuttle

Gail Chapline

Intercommunication Comes of Age—The Digital Age

John Hirko

Restoring Integration and Systems Thinking in a Complex Midlife Program

John Muratore

Electromagnetic Compatibility for the Space Shuttle

Robert Scully

Process Control

Glen Curtis

Steven Sullivan

David Wood

Alliant Techsystems, Inc. and United Space Alliance

Holly Lamb

Dennis Moore

David Wood

Michoud Assembly Facility

Jeffery Pilet

Kenneth Welzyn

Patrick Whipps

Pratt & Whitney Rocketdyne Manufacturing

Eric Gardze

Rockwell International and The Boeing Company

Bob Kahl

Larry Kauffman

NASA and the Environment—Compatibility, Safety, and Efficiency

Samantha Manning

Protecting Birds and the Shuttle

Stephen Payne

All complex systems require systems engineering that integrates across the subsystems to meet mission requirements. This interdisciplinary field of engineering traditionally focuses on the development and organization of complex systems. However, NASA applied systems engineering throughout the life cycle of the Space Shuttle Program—from concept development, to production, to operation and retirement. It may be surprising to many that systems engineering is not only the technical integration of complex space systems; it also includes ground support and environmental considerations. Engineers require the aid of many tools to collect information, store data, and interpret interactions between shuttle systems. One of the shuttle's legacies was the success of its systems engineering. Not only did the shuttle do what it was supposed to do, it went well beyond meeting basic requirements.

This section is about systems engineering innovations, testing, approaches, and tools that NASA implemented for the shuttle. Companies that developed, built, and maintained major shuttle components are highlighted. As manufacturers, contractors, NASA, and industry employees and management came and went, the shuttle stayed the same during its lifetime, primarily because of its well-honed process controls. All of these systems engineering advances are a legacy for the International Space Station and for future space vehicles.



Systems Engineering During Development of the Shuttle

Systems engineering is a complex, multilevel process that involves deconstructing a customers' overall needs into functions that the system must satisfy. But even in ordinary situations, that's just the beginning. Functional requirements are then allocated to specific components in the system. Allocated functions are translated into performance requirements and combined with design constraints to form requirements that a design team must satisfy. Requirements are then synthesized by a team of engineers into one or more concepts, which are traded off against each other. These design concepts are expanded into preliminary and detailed designs interspersed with reviews. Specialists from many disciplines work as a team to obtain a solution that meets the needs and requirements. Selected designs are translated into manufacturing, planning, procurement, operations, and program completion documents and artifacts.

Systems engineering for the Space Shuttle presented an extraordinary situation. The shuttle was the most

complex space vehicle for its time and, therefore, required the evolution of systems engineering with significantly advanced new tools and modeling techniques. Not only was the vehicle sophisticated, it required the expertise of many people. Four prime contractors and thousands of subcontractors and suppliers, spread across the United States, designed and built the major elements of the shuttle. The complexity of the element interfaces meant the integration of elements would present a major systems engineering challenge. One prime contractor was in charge of building the main engines, which were mounted inside the Orbiter. A different prime contractor built the Orbiter. A third prime contractor built the External Tanks, which contained the fuel for the main engines. And, a fourth prime contractor built the Solid Rocket Boosters. As problems occurred, they involved multiple NASA engineering organizations, industry partners, subject matter experts, universities, and other government agencies. NASA's ability to bring together a wide group of technical experts to focus on problems was extremely important. Thus, one legacy of the Space Shuttle was the success of its systems engineering. Not only did the shuttle do what it was

supposed to do, it went well beyond meeting basic requirements.

A discussion of all the systems engineering models and new tools developed during the lifetime of the Space Shuttle Program would require volumes. All elements of the Space Shuttle Program had successes and failures. A few of the most notable successes and failures in systems engineering are discussed here.

Change and Uncertainty

Space Shuttle Main Engines

NASA recognized that advancements were needed in rocket engine technology to meet the design performance requirements of the shuttle. Thus, its main engine was the first contract awarded.

A high chamber pressure combined with the amplification effect of the staged combustion cycle made this engine a quantum leap in rocket engine technology for its time. The engine also had to meet the multiple interface requirements to the vehicle, extensive operation requirements, and several design criteria. A major challenge for systems engineering was


Intercommunication Comes of Age—The Digital Age

As the shuttle progressed, it became evident that the existing communication system could not meet the multi-flow and parallel processing requirements of the shuttle. A new system based on digital technology was proposed and Operational Intercommunication System-Digital was born, and is now in its third generation. This system provided unlimited conferencing on 512 communication channels and support for thousands of end

users. The system used commercially available off-the-shelf components and custom-designed circuit boards.

Digital communication systems included, among other things, the voice communication system at Kennedy Space Center (KSC). The voice communication system needed to perform flawlessly 24/7, 365 days a year. This need was met by Operational Intercommunication

System-Digital—a one-of-a-kind communication system conceived, designed, built, and operated by NASA engineers and a team of support contractors. The system was installed in every major processing facility, office building, and various labs around KSC. This widespread distribution allowed personnel working on specific tasks to communicate with one another, even in separate facilities.



Dominic Antonelli

Commander, US Navy.

Pilot on STS-119 (2009) and STS-132 (2010).

"At the end of the day, people comprise the system that ultimately propelled the Space Shuttle Program to its stellar place in history. The future of space travel will forever be indebted to the dedication, hard work, and ingenuity of the men and women, in centers across the country, who transformed the dream into a tangible reality and established a foundation that will inspire generations to come."



that all of these requirements and design criteria were interrelated.

In most complex systems, verification testing is performed at various stages of the buildup and design. NASA followed this practice on previous vehicles. In component-level tests, engineers find problems and solve them before moving to the next higher assembly level of testing. The main engine components, however, were very large. Test facilities that could facilitate and perform the component and higher assembly level tests did not exist. The valves alone required a relatively large specialized test facility. Plans to build such facilities had been developed, but there was not enough time to complete their construction and maintain the schedule. Therefore, the completed main engine became the test bed.

A concurrent engineering development philosophy associated with the shuttle forced the engine to be its own test bed. The engine test stands at Stennis Space Center in Mississippi were already in place, so NASA decided to assemble the engines and use them as the breadboard or facility to test the components. This was a risky scenario. The engine proved to be unforgiving. NASA lost 13 engines from catastrophic failures on the test stand

before first flight. Each of these failures was a rich learning experience that significantly enabled the engineers to improve the engine's design. Still, at times it seemed the technical challenges were insurmountable.

Another philosophy that prevailed in the development of the main engines was "test, test, and test some more." Testing was key to the success of this shuttle component. Technicians conducted tests with cracked blades, rough bearings, and seals with built-in flaws to understand the limitations. By late 1979, as noted in a paper written by Robert Thompson, Space Shuttle manager at the time: "We have conducted 473 single engine tests and seven multiple engine tests with a cumulative total running time of 98 times mission duration and with 54 times mission duration at the engine rated power level. Significant engine test activities still remain and must be completed successfully before the first flight, but the maturity of this vital system is steadily improving."

The test, test, and test some more philosophy reduced risk, built robustness, and added system redundancy. Testing also allowed engineers to understand interactions of failures with other systems during

the 30 years of the program. In all, the main engines were upgraded three times. These upgrades improved the engines' performance and reliability, reduced turnaround costs, and were well-planned system engineering efforts.

Throughout the life of the Space Shuttle Program—and through many technical challenges and requirement changes—the main engine not only performed, but was also a technological leap for spacecraft rocket engines.

Where Was Systems Engineering When the Shuttle Needed It Most?

Thermal Protection System

Early development problems with the Orbiter's Thermal Protection System probably could have been avoided had a systems engineering approach been implemented earlier and more effectively.

The Thermal Protection System of the Orbiter was supposed to provide for the thermal protection of the structure while maintaining structural integrity. The engineers did a magnificent job in designing tiles that accepted, stored, and dissipated the heat. They also created a system that maintained the aerodynamic configuration. However, early in the process, these engineers neglected to design a system that could accept the loads and retain the strength of the tiles. Furthermore, it was not until late in the Thermal Protection System development process that NASA discovered a major problem with the attachment of tiles to the Orbiter's aluminum skin surfaces.

In 1979, when Columbia—the first flight Orbiter—was being ferried from Dryden Flight Research Center in California to Kennedy Space Center in Florida on the back of the 747 Shuttle Carrier Aircraft, several tiles fell off. This incident focused NASA's



attention on the tile attachment problem. The solution ultimately delayed the maiden flight of Columbia (Space Transportation System [STS]-1) by nearly 1½ years. The problem resided in the bond strength of the tiles, which was even lower than the overall low strength of the tile material. Tile load analyses kept showing increasing loads and lower margins on tile strength. This low bond strength was related to stress concentrations at the bondline interface between the tile and the strain isolation pad. Attachment of the tiles to the Orbiter's aluminum skin required that the strains from structural deflections be isolated from the tiles. In other words, the tiles could not be bonded directly to the Orbiter structure.

Strain isolation was accomplished with Nomex® felt pads bonded to the structure. In turn, the tiles bonded to the pads. Needling of the Nomex® pads through the thickness to control thickness resulted in straight through fibers ("stiff spots") that induced point loads in the bottom of the tiles. These point loads caused early localized failure of the tile material at the bondline. This did not meet design requirements.

After more than 1 year of intense, around-the-clock proof testing, bonding, removing, and re-bonding of tiles on the vehicle at Kennedy Space Center, tile densification proved to be the solution. Stress concentrations from the strain isolation pad were smoothed out and the full tile strength was regained by infusing the bottom of the tiles, prior to bonding, with a silica-based solution that filled the pores between tile fibers for a short distance into the bottom of the tile. This example demonstrates that a systems approach to the tile design, taking into consideration not only the thermal performance of the tile but also the structural integrity, would have allowed the tile attachment problem to be solved earlier in the design process.

The Importance of Organizational Structure

The structure of the Space Shuttle Program Systems Integration Office was a key element in the successful execution of systems engineering. It brought together all shuttle interfaces and technical issues. Design and performance issues were brought forward there. The office, which integrated all technical disciplines, also had a technical panel structure that worked the technical details from day to day.

The panels were composed of engineers from multiple NASA centers, prime contractors, and subcontractors.

NASA also brought in technical experts when needed.

These panels varied in size. The frequency of discussions depended on the technical areas of responsibility and the difficulty of the problems encountered. The panels operated in an environment of healthy tension, allowing for needed technical interchange, questioning, and probing of technical issues. The technical panel structure has been recognized as a significant and an effective means to manage complex systems.

Initially, there were 44 formalized panels, subpanels, and working groups in the Space Shuttle Program Office.





However, because of the complexity, by 1977 the number had grown to 53 panels, subpanels, and working groups. These critical reviews provided guidance to maintain effective and productive technical decisions during the shuttle development phase. Also during this phase of the program, NASA established the definition and verification of the interfaces and associated documentation, including hazard analysis and configuration control.

Biggest Asset— People Working Together

Owen Morris, manager of the Systems Integration Office from 1974 to 1980, was an effective and a respected manager. When asked to describe the biggest challenge of that position, Owen answered, “People. Of course, all the people involved had their own responsibilities for their part of the program, and trying to get the overall program put together in the most efficient manner involved people frequently giving up part of their capability, part of their prerogative, to help a different part of the program, solve a problem, and do it in a manner that was better for everyone except them. And, that’s a little difficult to convince people to do that. So, working with people, working with organizations, and getting them to work together in a harmonious manner was probably the most difficult part of that.”

The challenge of getting people to work together successfully has been an enduring one. NASA stepped up to multiple challenges, including that of having various people and organizations working together toward a common goal. By working together, the space agency engineered many successes that will benefit future generations.

Restoring Integration and Systems Thinking in a Midlife Program

Aviation lore says that, during World War II, a heavily overworked crew chief confronted an aircraft full of battle damage and complained, “That’s not an airplane, that’s a bunch of parts flying in loose formation.”

One of the greatest challenges during system development is transforming parts into a fully integrated vehicle. Glenn Bugos’ book titled *Engineering the F-4 Phantom II* is subtitled *Parts into Systems* in recognition of this challenge. NASA also long realized this. In the standard NASA cost model for space systems, the agency planned that 25% of a program’s development effort would go into systems engineering and integration. Efforts made during the initial development of the shuttle to ensure its integrated performance led to a successful and an enduring design.

NASA Learns an Expensive Lesson

NASA’s experience in human spacecraft prior to the shuttle was with relatively short-lived systems. The agency developed four generations of human spacecraft—Mercury, Gemini, Apollo, and Skylab—in fewer than 15 years. Designers and project managers intuitively anticipated rapid replacement of human space systems because, at the time of shuttle development, they had no experience to the contrary. The initial design parameters for the Orbiter included 100 missions per Orbiter in 10 years. During the design phase, NASA did not plan for the 30-year operational life the shuttle actually flew.

The space agency, therefore, had no experience regarding the role of systems engineering and integration during the extended operational part of a system life cycle. Given the cost of a strong systems engineering and integration function, this was a topic of significant debate within NASA, particularly as budgets were reduced. As late as 1990—9 years after the shuttle’s first flight—the systems engineering and integration effort was approximately \$160 million per year, or approximately 6.4% of the \$2.5 billion shuttle annual budget. Starting in 1992, to meet reduced operating budgets, this level of resource came under scrutiny. It was argued that, given major development of the shuttle system was complete, all system changes were under tight configuration control and all elements understood their interfaces to other elements, the same level of systems engineering and integration was no longer required. The effort was reduced to 2.2% of the shuttle annual budget in 1992. Occurrences of in-flight anomalies were decreasing during this period, thereby lending to the belief that the proper amount of integration was taking place.

This seemed to be a highly efficient approach to the problem until the loss of Columbia in 2003. In retrospect, the Columbia Accident Investigation Board determined there were clear indicators that the program was slowly losing the necessary degree of systems engineering and integration prior to the loss of Columbia. Critical integration documentation no longer reflected the vehicle configuration being flown. Furthermore, the occurrence of integrated anomalies was increasing over the years.



Crucial Role of Systems Engineering

Known Changes

Change was constantly occurring in the shuttle systems. Changes with known effects required a large and expensive integrated engineering effort but were usually the easiest to deal with. For example, when NASA upgraded the Space Shuttle Main Engines to a more-powerful configuration, a number of changes occurred in terms of avionics, electrical, and thrust performance. These changes had to be accommodated by the other parts of the system.

Known changes with unknown effects were more difficult to deal with. For example, as a cost-reduction effort, NASA decided not to replace the connectors on the Orbiter umbilicals after every flight. At the time, NASA did not know that the Solid Rocket Booster exhaust and salt-spray environment of the pad created corrosion on the connectors. This corrosion would eventually interrupt safety-critical circuits. On Space Transportation System (STS)-112 (2002), half the critical pyrotechnic systems, which release the shuttle from the launch pad, did not work. Because the systems had redundancy, the flight launched successfully.

Unknown Changes—Manufacturing Specification

There were many sources of unknown change during the Space Shuttle Program. First, the external environment was continually changing. For example, the electromagnetic environment changed as radio-frequency sources appeared and disappeared in terrain over which the shuttle flew. These sources could influence the performance of shuttle systems.

Second, the characteristics of new production runs of materials such as adhesives, metals, and electronic components changed over time. It was impossible to fully specify all characteristics of all materials on a large system. Changes in assembly tooling or operators could have resulted in a product with slightly different characteristics. For instance, major problems with fuel quality circuits caused launch delays for flights after the Columbia accident. The circuits were intended to identify a low fuel level and initiate engine shutdown, thus preventing a probable engine catastrophe. These circuit failures were random. While these anomalies remained unexplained, the circuit failures seemed to stop after improvements were made to the engine cutoff sensors. However, following another failure on STS-122 (2008), the problem was isolated to an electrical connector on the hydrogen tank and was determined to be an open circuit at the electrical connector's pin-to-socket interface. The increased failure rate was likely caused by a subtle change to the socket design by the vendor, combined with material aging within the connector assembly. The connector was redesigned, requiring soldering the sockets directly to the pins.

Solution—Systems Engineering


The only way to deal with known and unknown change was to have a significant effort in systems engineering and integration that monitored integrated flight performance and was attuned to the issues that could impact a system. One of the best approaches for maintaining this vigilance was comparing in-flight anomalies to established analyses of hazards to the integrated system. These integrated hazard analyses were produced at the start of the program but had not been updated at the time of the Columbia

accident to reflect the present vehicle configuration. Further, the in-flight anomaly process was not tied to these analyses. In the period before Return to Flight, the systems engineering and integration organization tried to fix these analyses but determined the analyses were so badly out of date that they had to be completely redone. Thus, systems engineering and integration replaced 42 integrated hazards with 35 new analyses that used fault-tree techniques to determine potential causes of hazards to the integrated system. These analyses were also tied into a revamped in-flight anomaly process. Any problem occurring in flight that could cause a hazard to the integrated system required resolution prior to the next flight.

Preparing for Return to Flight After the Columbia Accident

When internal NASA evaluations and the Columbia Accident Investigation Board determined that shuttle systems engineering and integration would need to be rebuilt, NASA immediately recognized that systems engineering and integration could not be rebuilt to 1992 levels. There were simply not enough available, qualified systems engineers who were familiar with the shuttle configuration. Further, it was unlikely that NASA could afford to maintain the necessary level of staffing. NASA accomplished a modest increase of about 300 engineers by selective hiring. Also, NASA worked with the Aerospace Corporation (California), along with establishing agreements with other NASA centers, such as integration personnel at Marshall Space Flight Center and Kennedy Space Center. This returned systems engineering and integration activities to 1995 levels. More impressive was the way in which these resources were deployed.

The most immediate job for systems engineering and integration during this



Left photo: Ames Research Center wind tunnel test.

Right photo: Aerothermal test at Calspan-University of Buffalo Research Center.

period was determining design environments for all redesigns mandated by the Columbia Accident Investigation Board. The standard techniques for establishing design environments prior to this effort involved constructing environment changes to the basic environments by making conservative calculations based on the nature of the change.

A large number of configuration changes over the years resulted in an accumulation of conservative design environments. However, this cumulative approach was the only basis for estimating the environments. A new baseline effort would have required extensive calculations and ground tests. For the Return to Flight effort, systems engineering and integration decided to re-baseline the critical design environments to eliminate non-credible results. Fortunately, technology had advanced significantly since the original baseline environments were constructed in the 1970s. These advances enabled greater accuracy in less time.

The shuttle aerodynamics model was refurbished to the latest configuration for aerodynamics and aerodynamic loads. Shuttle wind tunnel tests were completed at Ames Research Center in California and the Arnold Engineering Development Center in Tennessee.

Engineers employed new techniques, such as pressure-sensitive paint and laser velocimetry in addition to more advanced pressure and force instrumentation. The purpose of these tests was to validate computational fluid dynamics models because design modifications were evolving as the design environments were being generated. Thus, continued wind tunnel tests could not generate the final design environments. Validated computational fluid dynamics models were necessary to generate such environments for the remainder of the Space Shuttle Program to avoid the accumulation of conservative environments.

Engineers performed similar tests using the aerothermal model at the Calspan-University of Buffalo Research Center (New York) shock tunnel. Engineers used a combination of computational fluid dynamics and other engineering methods to generate an updated thermal database.

Another major task for systems engineering and integration was to understand the debris transport problem. A 0.76-kg (1.67-pound) piece of foam debris was liberated from the External Tank. This foam debris was responsible for the damage that caused the Columbia accident. Systems engineering and integration enabled

engineers to identify the transport paths of debris to the shuttle to determine the hazard level of each debris item as well as determine the impact velocities that the structure would have to withstand. When analysis or testing revealed the elements could not withstand impact, systems engineering and integration worked with the debris-generating element to better understand the mechanisms, refine the estimated impact conditions, and determine whether debris-reduction redesign activities were sufficient to eliminate or reduce the risk. To understand debris transport, NASA modeled the flow fields with computational fluid dynamics and flight simulation models. Fortunately, NASA had entered into an agreement, post-Columbia, to create the world's largest supercomputer at Ames Research Center. This 10,240-element supercomputer came on line in time to perform extensive computational fluid dynamics and simulation analysis of debris transport.

Debris Transport During Launch Remained a Potential Hazard

NASA cataloged both the size and the shape of the debris population as well as the debris aerodynamics over a wide speed range. A large part of this



NASA validated computational fluid dynamics and flight simulation models of the foam debris in flight tests using the Dryden Flight Research Center (California) F-15B Research test bed aircraft. In these tests, debris fell from foam panels at simulated shuttle flight conditions. High-speed video cameras captured the initial flight of the foam divots.



effort involved modeling the flight characteristics of foam divots that came off of the tank. NASA first addressed this problem by firing small plastic models of foam divot shapes at the NASA Ames Research Center, California, ballistic range. When these results correlated well with computational fluid dynamics, the agency conducted more extensive tests. Engineers tested flight characteristics of foam debris in the Calspan-University of Buffalo Research Center tunnel and Dryden Flight Research Center, California. Results showed that foam would stay intact at speeds up to Mach 4 and, therefore, remain a potential hazard.

Other Return to Flight Activities

Two other major tasks were part of the systems engineering and integration Return to Flight effort. The first task involved integrated test planning to ensure that the system design was recertified for flight. The second task was to install additional instrumentation and imagery acquisition equipment to validate the performance of system design changes.

The diversity of integrated system testing was remarkable. Integrated tests included the first-ever electromagnetic interference tests run on the shuttle system. NASA ran a test to determine the effects of the crawler transporter on the vibration/fatigue of shuttle structures. This effort required construction of improved integrated structural models. First performed on a limited scale during the Return to Flight period, this effort expanded under Marshall Space Flight Center leadership. The integrated test effort also included two full-up tanking tests of the shuttle system. In addition to validating the performance of the new foam system on the tank, these tanking

tests discovered two major problems in the shuttle: failures of the propellant pressurization system and problems with the engine cutoff sensors.

The instrumentation added to the shuttle system as part of the systems engineering and integration effort was also diverse. NASA added instrumentation to the External Tank to understand the vibration and loads on major components attached to the skin. These data proved vital after Return to Flight assessment because a loss of foam associated with these components required additional modification. This instrumentation gave the program the confidence to make these modifications. NASA also added instrumentation to help them understand over-pressure effects on the shuttle due to ignition transients of the Space Shuttle Main Engine and motion of the Orbiter-ground system umbilicals. The agency added ground-based radar and video imaging equipment to provide greater visibility into the debris environment and validate design modifications.

Integration Becomes the Standard

NASA learned some difficult yet valuable lessons about the importance of systems engineering and integration over the course of the Space Shuttle Program—especially in the years following the loss of Columbia. The lack of systems engineering and integration was a contributing cause to the accident. The shuttle had become “a collection of parts flying in loose formation.” It took a major engineering effort over a 2-year period to reestablish the proper amount of integration. This effort significantly improved the shuttle system and laid the groundwork and understanding necessary for the successful flights that followed.

Electromagnetic Compatibility for the Space Shuttle

Electromagnetic compatibility is extremely complex and far reaching. It affects all major vehicle engineering disciplines involving multiple systems and subsystems and the interactions between them. By definition, electromagnetic compatibility is the capability of electrical and electronic systems, equipment, and devices to operate in their intended electromagnetic environment within a defined margin of safety, and at design levels of performance. But, that is just the beginning. This must be accomplished without causing unacceptable degradation as a result of any conducted or radiated electromagnetic energy that interrupts, obstructs, or otherwise limits the effective performance of telecommunications or other electrical and electronic equipment.

Design and Verification Requirements— A Learning Process

In 1973—when NASA was first defining the shuttle systems—military models offered the best available means of providing control of the system design leading to acceptable levels of electromagnetic compatibility. Previous requirements for Mercury, Gemini, and Apollo were cut from the same cloth, but none of those programs had a vehicle that could compare to the shuttle in terms of size and complexity.

Admittedly, these comprehensive requirements addressed a multiplicity of concerns. These included: subsystem criticality; degradation criteria; interference and susceptibility control;



wiring and cable design and installation; electrical power; electrical bonding and grounding; control of static electricity and its effects; electromagnetic hazards to personnel, explosives, and ordnance; and definition of, and design for, the external electromagnetic environment.

Detailed design and verification requirements for protection from the damaging effects of lightning were also included and developed independently by NASA. These shuttle lightning requirements became the foundation for a plethora of military and commercial aerospace requirements, culminating in a detailed series of Society of Automotive Engineers documents universally employed on an international basis.

A Custom Fit Was Needed

Unfortunately, without a solid basis for the tailoring of requirements, shuttle electromagnetic compatibility engineers chose to levy the baseline requirements with virtually no change from previous Apollo efforts. Although this was a prudent and conservative approach, it led to misinterpretation and misapplication of many requirements to the shuttle. As a result, NASA granted an unacceptably large number of waivers for failure to comply with the requirements. The problem continued to grow until 2000, at which time NASA made a major effort to completely review and revise the electromagnetic compatibility requirements and compliance approach. This effort eliminated or tailored requirements so that the content was directly and unequivocally applicable to the shuttle. This effort also allowed for a systematic and detailed revisitation of previously granted waivers against the backdrop of the new requirements' definitions.

Making Necessary Adjustments...and Succeeding

Original requirements and new requirements were tabulated together to facilitate direct comparison. For each set of requirements, NASA needed to examine several characteristics, including frequency range, measurement circuit configuration, test equipment application, and the measured parameter limits. As an example, certain conducted emissions requirements in the original set of requirements measured noise currents flowing on power lines whereas the equivalent new requirements measured noise voltages on the same power lines. To compare limits, it was necessary to convert the current limits to voltage limits using the linear relationship between voltage, current, and circuit impedance.

In other cases, frequency bandwidths used for testing were different, so NASA had to adjust the limits to account for the bandwidth differences.

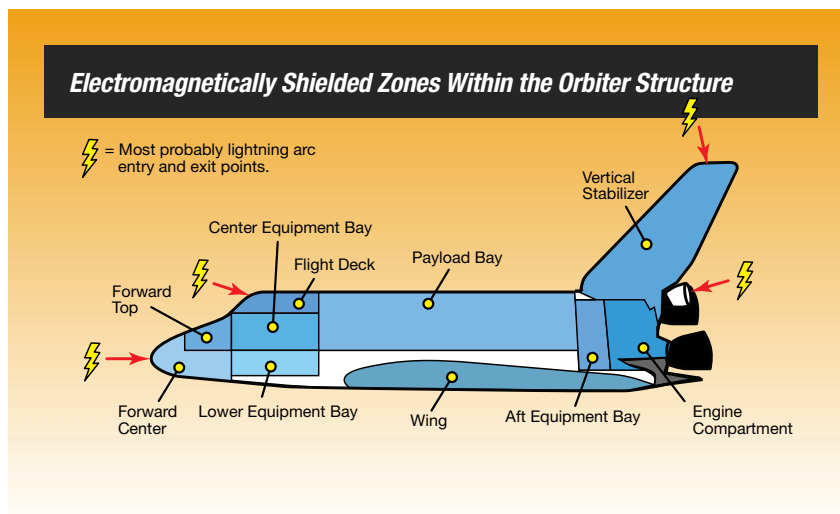
In all, NASA engineers were able to work through the complexity of electromagnetic compatibility—to follow all of the threads inherent in the vehicle's multiple systems and subsystems—and find a way to tailor the requirements to accommodate the shuttle.

Process Control

The design and fabrication of the Space Shuttle's main components took place in the early 1970s while Richard Nixon was president. The Space Shuttle was assembled from more than 2.5 million parts that had to perform per design with very little margin of error. NASA constantly analyzed and refurbished flight systems and their components to ensure performance. The success of the Space Shuttle Program was due in great part to diligent process control efforts by manufacturing teams, contractors, and civil service engineers who carefully maintained flight hardware.

Five Key Elements Ensure Successful Process Control

Process control consists of the systems and tools used to ensure that processes are well-defined, perform correctly, and are maintained such that the completed product conforms to requirements. Process control managed risk to ensure safety and reliability in a complex system. Strict process control practices helped prevent deviations that could have caused or contributed to incidents, accidents, mishaps, nonconformances, and in-flight





anomalies. As defined by NASA, the five key elements of a process are: people, methods/instructions, materials, equipment, and environment. It has been long understood that qualified, conscientious people are the heart of any successful operation. High-quality process control efforts require skilled, detail-oriented individuals who understand and respect the importance of process and change control. The methods or instructions of a process, often called “specifications” or

“requirements,” are those documented techniques used to define and perform a specific process. The term “equipment” refers to the tools, fixtures, and facilities required to make products that meet specifications and requirements while “materials” refers to both product and process materials used to manufacture and test products. Finally, the environmental conditions required to properly manufacture and test products must also be maintained to established standards to ensure safety and reliability.

Solid Engineering Design— A Fundamental Requirement

A clear understanding of the engineering design is fundamental when changes occur later in a program’s life. Thousands of configuration changes occurred within the Space Shuttle Program. These changes could not have been made safely without proper process controls that included a formal configuration control system. This

Alliant Techsystems, Inc. and United Space Alliance


The signature twin reusable solid rocket motors of the Space Shuttle carried the fingerprints of thousands of people who designed, manufactured, tested, and evaluated the performance of these workhorse motors since 1982. The manufacturing facility in Promotory, Utah, is now owned and operated by Alliant Techsystems, Inc. (ATK). Originally developed to manufacture and test large-scale rocket motors for intercontinental ballistic missiles, the site provided 72% of the liftoff thrust to loft each shuttle beyond Earth’s bounds.

The Assembly Refurbishment Facility complex—managed and operated by United Space Alliance (USA), headquartered in Houston, Texas—is located at Kennedy Space Center, Florida. The complex began operations in 1986 and was the primary integration and checkout facility for boosters. Refurbished and new hardware were assembled and submitted to rigorous



Technicians process the solid rocket motor case segments at the ATK case lining facility in Utah.

testing to assure the assemblies were ready for human-rated flight. The facility was equipped to handle assembly, testing, and troubleshooting of thrust vector control systems, avionics, and recovery systems for the Space Shuttle Program.



Solid Rocket Booster case preparation.



Propellant mixing.



Solid Rocket Booster aft skirt processing at the Assembly and Refurbishment Facility at Kennedy Space Center.

© Pratt & Whitney RocketDyne. All rights reserved.



Michoud Assembly Facility

By the end of the Space Shuttle Program, NASA's Michoud Assembly Facility—located near New Orleans, Louisiana, and managed by Marshall Space Flight Center in Huntsville, Alabama—delivered 134 External Tanks (ETs) for flight. Two additional tanks were built but not scheduled to fly, and three assemblies were delivered for major tests, resulting in a total of 139 tanks. As one of the world's largest manufacturing plants, Michoud's main production building measured 17 hectares (43 acres) under one roof, including a 61-m (200-ft) vertical assembly building, and a port that permitted transportation of ETs via oceangoing barges and towing vessels to Kennedy Space Center in Florida.

ETs were produced at Michoud by prime contractor Lockheed Martin (headquartered in Bethesda, Maryland) over a 37-year period. The contractor procured parts and materials from hundreds of subcontractors across the country. In full production, 12 tanks were in various phases of production across the facility—each tank requiring approximately 3 years to complete. Each ET included over 0.8 km (0.5 miles) of welds, thousands of rivets and bolts, redundant inspections within each process, and sophisticated pressure and electrical testing.

Throughout the history of the program, Michoud continually improved the processing, materials, and components of ETs. Improvements included the introduction of a stronger, lighter aluminum-lithium alloy—which saved over 2.7 metric tons (3 tons) of weight—and transitioning to virtually defect-free friction stir welding. Additionally, Michoud developed thermal protection foam spray systems and process controls that reduced weight and minimized foam loss during the extreme environments of flight.



Liquid oxygen tank.



Liquid oxygen tank and intertank in a checkout cell.



Liquid hydrogen tank showing slosh and vortex baffle inside.



External Tank processing.

system involved the use of review boards, material review analyses, and tool controls.

A Team Effort

Hardware for the Space Shuttle Program was manufactured by a broad supplier base using a variety of processes. If these processes were not controlled, a deterioration of the end product could have occurred, thereby increasing risk. In essence, NASA depended on the process controls at over 3,000 flight hardware suppliers' facilities across the United States. Any subtle changes or deviations from any established processes could have negatively affected the outcome.

Think of the thousands of vendors and processes that might have affected manufacturing—from material pedigree to the material of gloves worn by a technician. All of these nuances affected the outcome of the product. Coordination and communication between NASA and its manufacturers were critical in this complicated web of hardware suppliers. The Space Shuttle was only as strong as its weakest link.

Strong process controls resulted in highly predictable processes. Built-in tests were critical because many flight components/systems could not be tested prior to their actual use in flight. For example, Thermal Protection Systems, pyrotechnics, and solid rocket motors could only be tested at the manufacturer's facilities before they were installed aboard the shuttle. This fact demonstrated once again that NASA was highly dependent on the integrity of its hardware suppliers to follow the tried and true "recipe" of requirements, materials, people, and processes to yield predictable and reliable components.

Processes Continue Well Beyond Flight

Because shuttles were reusable vehicles, process control was also vital to refurbishment and postflight evaluation efforts. After each flight, NASA closely monitored the entire vehicle to evaluate factors such as heat exposure, aging effects, flight loads, shock loads, saltwater intrusion, and other similar environmental impacts. For example, did you know that each heat tile that protected the underbelly of the vehicle from the extreme heat of re-entry into Earth's atmosphere was numbered and checked following each flight? Tiles that did not pass inspection were either repaired or replaced. This effort was a major undertaking since there were 23,000 thermal protection tiles.

Postflight recovery and inspections were an important part of process control. For example, NASA recovered the Solid Rocket Boosters, which separated from the vehicle during launch and splashdown in the Atlantic Ocean, and brought them back to Kennedy Space Center in Florida where they were examined and inspected. These standardized forensic inspections provided valuable data that determined whether the booster system operated within its requirements and specifications. Data collected by the manufacturer represented the single most important feedback process since this system had to function as intended every time without the ability to pretest.

Best Practices Are Standard Practice

Each of NASA's manufacturers and suppliers had unique systems for process control that guaranteed the integrity of the shuttle's hardware.

Pratt & Whitney Rocketdyne Manufacturing

The Space Shuttle Main Engine required manufacturing and maintenance across the entire United States. Pratt & Whitney Rocketdyne (Canoga Park, California), under contract to NASA, developed the main engine, which successfully met the challenges of reusability, high performance, and human-rated reliability. With every launch, the team continued to make improvements to render it safer and more reliable.



High-pressure fuel turbopump recycling.

The Pratt & Whitney Rocketdyne facility at the West Palm Beach, Florida, campus designed and assembled the critical high-pressure turbomachinery for the shuttle. The high pressures generated by these components allowed the main engine to attain its extremely high efficiency. At the main facility in Canoga Park, California, the company fabricated and assembled the remaining major components. The factory included special plating tanks for making the main combustion chamber (the key components to attain high thrust with the associated high heat transfer requirements), powerhead (the complex structural heart of the engine), and nozzle (another key complex component able to withstand temperatures of 3,300°C [6,000°F] degrees during operation). In addition, the company employed personnel in Huntsville, Alabama, and Stennis Space Center in Mississippi. The Huntsville team created and tested critical software. The Stennis team performed testing and checkout of engines and engine components before delivery to the launch site. Finally, at Kennedy Space Center in Florida, Pratt & Whitney Rocketdyne personnel performed all the hands-on work required to support launch, landing, and turnaround activities.



Space Shuttle Main Engine assembly.



Rockwell International and The Boeing Company

Rockwell of Downey, California (now Boeing) executed the Orbiter design, development, test, and evaluation contract, the production contract, and the system integration contract for the mated shuttle vehicle. Engineers were the primary producers of specifications, vehicle loads/environments, analysis, drawing release, certification/qualification testing, and certification documentation. Engineers performed key system-level integration and testing for many Orbiter subsystems including software, avionics hardware, flight controls/hydraulics, and


thermal protection. At this same location, technicians manufactured the crew module, forward fuselage, and aft fuselage, which were integrated into the Orbiter at the Boeing facility in Palmdale, California.

Boeing engineers, technicians, and support personnel assembled and tested all six Space Shuttle Orbiter vehicles. The first shuttle vehicle, Enterprise, was delivered in January 1977. Being a non-orbital vehicle, it was used for fit checks, support equipment procedures, and the Approach and Landing Test Program conducted at

Dryden Flight Research Center on the Edwards Air Force Base runway in California beginning in 1977. Columbia, the first space-rated Orbiter, was delivered in the spring of 1979 and later flew the Space Shuttle Program's maiden voyage in April 1981. Challenger was rolled out in 1982, followed by Discovery in 1983 and Atlantis in 1985. The newest shuttle, Endeavour, was authorized following the loss of Challenger in 1986 and was delivered in April 1991. From 1985 to 2001, engineers performed eight major modifications on the Orbiter fleet.



Orbiter assembly.



© The Boeing Company. All rights reserved.

The communication and establishment of specific best practices as standards helped the program improve safety and reliability over the years. The following standards were the minimum process control requirements for all contractors within the Space Shuttle Program:

- Detect and eliminate process variability and uncoordinated changes.
- Eliminate creep—or changes that occur over time—through process controls and audits.
- Understand and reduce process risks.
- Identify key design and manufacturing characteristics and share lessons learned that relate to the processes.

- Be personally accountable and perform to written procedures.
- Promote process control awareness.
- Identify and evaluate changes to equipment and environment.
- Capture and maintain process knowledge and skills.

NASA witnessed a significant evolution in their overall process control measures during the shuttle period. This lengthy evolution of process control, a continuous effort on the part of both NASA and its contractors, included multiple initiatives such as:

- establishing reliable processes
- monitoring processes
- reinforcing the process-control philosophy or “culture”
- maintaining healthy systems

Establishing reliable processes included open communications (during and after the design process) among numerous review boards and change boards whose decisions dictated process-control measures. Monitoring processes involved postflight inspections, safety management systems, chemical fingerprinting, witness panels, and other monitoring procedures. Process



control also referred to relatively new programs like the “Stamp and Signature Warranty” Program where annual audits were performed to verify the integrity of products/components for the shuttle era. Finally, maintaining healthy systems focused on sustaining engineering where design or operating changes were made or corrective actions were taken to enhance the overall “health” of the program.

An Enduring Success

Although NASA’s process control measures have always been rigorous, additional enhancements for improved communication and information-sharing between shuttle prime contractors and suppliers

created highly restrictive, world-class standards for process control across the program. Many of these communication enhancements were attainable simply because of advances in technology. The computer, for example, with its increased power and capabilities, provided faster and better documentation, communication, data tracking, archiving, lot number tracking, configuration control, and data storage. As manufacturers, contractors, and other businesses came and went—and as employees, managers, and directors came and went—the program stayed the same over its lifetime and continued to operate successfully primarily because of its well-honed process-control measures.

NASA and the Environment—Compatibility, Safety, and Efficiency

As conscientious stewards of US taxpayers dollars, NASA has done its part to mitigate any negative impacts on the wildlife and environment that the agency’s processes may impart. For NASA, it is not about technical issues; in this case, it is about the coexistence of technology, wildlife, and the environment.

Compatibility

The 56,700 hectares (140,000 acres) controlled by Kennedy Space Center (KSC) symbolize a mixture of technology and nature. Merritt Island National Wildlife Refuge was established in 1963 as an overlay of the center. The refuge consists of various habitats: coastal dunes; saltwater estuaries and marshes; freshwater impoundments; scrub, pine flatwoods; and hardwood hammocks. These areas provide habitat for more than 1,500 species of plants and animals. Hundreds of species of birds reside there year-round, with large flocks of migratory waterfowl arriving from the North and staying for the winter. Many endangered wildlife species are native to the area. Part of KSC’s coastal area was classified as a national seashore by agreement between the NASA and the Department of the Interior.

Most of the terrain is covered with extensive marshes and scrub vegetation, such as saw palmettos, cabbage palm, slash pine, and oaks. Citrus groves are in abundance, framed by long rows of protective Australian pine. More than 607 hectares (1,500 acres) of citrus groves are leased to individuals who

The Case of the Chloride Sponges

Let’s look at “The Case of the Chloride Sponges” to further demonstrate the importance of process control and the complexities of maintaining the Space Shuttle fleet. Postflight maintenance requirements included applying a corrosion inhibitor (sodium molybdate) to the Space Shuttle Main Engine nozzles. Following the STS-127 (2009) flight, engineers observed increased nozzle corrosion instances in spite of the application of the corrosion inhibitor. A root-cause investigation found that the sponges used to apply the corrosion inhibitor contained high levels of chlorides. Apparently, the sponges being used to apply the corrosion inhibitor were themselves causing more corrosion.

It was determined that the commercial vendor for the sponges had changed their sponge fabrication process. They began adding magnesium chloride for mold prevention during their packaging process and since NASA did not have a specification requirement for the chloride level in the sponges, the sponge fabrication change initially went unnoticed. To solve this problem, NASA added a requirement that only chloride-free sponges could be used. The agency also added a specification for alternate applicator/wipes. Case closed!





tend to the trees and harvest their fruit. Beekeepers maintain the health of the trees by collecting honey from—and maintaining—the hives of bees essential to the pollination of the citrus trees. Merritt Island National Wildlife Refuge manages the leases. Other NASA centers such as White Sands Test Facility and Wallops Flight Facility are also close to National Wildlife Refuges.

Safety

There is a limit as to what NASA can do to actually protect itself from the wildlife. During launch countdown of Space Transportation System (STS)-70 on Memorial Day 1995, the launch team discovered a pair of northern flicker woodpeckers trying to burrow a nesting hole in the spray-on foam insulation of the shuttle External Tank on Pad B. Spray-on foam insulation was comparable to the birds' usual nesting places, which include the soft wood of palm trees or dead trees. However, on reaching the aluminum skin of the tank beneath the spray-on foam insulation layer, the woodpeckers would move to a different spot on the tank and try again. In the end, there were at least 71 holes on the nose of the tank that couldn't be repaired at the pad. As a result, the stack was rolled back to the Vehicle Assembly Building for repairs to the damaged insulation.

The problem of keeping the woodpeckers from returning and continuing to do damage to the tank's spray-on foam insulation proved to be complex. The northern flicker is a protected species so the birds could not be harmed. In NASA fashion, shuttle management formed the Bird Investigation Review and Deterrent (BIRD) team to research the flicker problem and formulate a plan for keeping the birds away from the pads.

After studying flicker behavior and consulting ornithologists and wildlife experts, the team devised a three-phase plan. Phase 1 of the plan consisted of an aggressive habitat management program to make the pads more unattractive to flickers and disperse the resident population of these birds. NASA removed palm trees, old telephone poles, and dead trees from the area around the pads. The agency allowed the grass around the pad to grow long to hide ants and other insects—the flickers' favorite food. Phase 2 implemented scare and deterrent tactics at the pads. NASA used plastic owls, water sprays, and "scary eye" balloons to make the area inhospitable to the birds and frighten them away without injuring them. Phase 3 involved the implementation of bird sighting response procedures. With the BIRD team plans in place and the flickers successfully relocated, STS-70 was able to launch approximately 6 weeks later.

Woodpeckers are not the only form of wildlife attracted to the External Tank. On STS-119 (2009), a bat was found clinging to Discovery's external fuel tank during countdown. Based on images and video, a wildlife expert said the small creature was a free tail bat that likely had a broken left wing and some problem with its right shoulder or wrist. Nevertheless, the bat stayed in place and was seen changing positions from time to time. The temperature never dropped below 15.6°C (60°F) at that part of the tank, and infrared cameras showed that the bat was 21°C (70°F) through launch. Analysts concluded that the bat remained with the spacecraft as it cleared the tower. This was not the first bat to land on a shuttle during a countdown. Previously, one landed on the tank during the countdown of STS-90 (1998).

Another species that NASA dealt with over the life of the Space Shuttle Program was a type of wasp called a mud dauber. Although the mud daubers aren't very aggressive and don't pose an immediate threat to people, the nests they build can pose a problem. Mud daubers tend to build nests in small openings and tubes such as test ports. This can be an annoyance in some cases, or much more serious if the nests are built in the openings for the pitot-static system (i.e., a system of pressure-sensitive instruments) of an aircraft. Nests built in these openings can affect functionality of the altimeter and airspeed indicator.

Efficiency

In keeping with imparting minimal negative impact on the environment, NASA also took proactive steps to reduce energy usage and become more "green." At KSC, NASA contracted several multimillion-dollar energy projects with Florida Power & Light Company that were third-party-financed projects. There was no out-of-pocket expense to NASA. The utility was repaid through energy savings each month. The projects included lighting retrofits; chilled water modifications for increased heating, ventilation, and air-conditioning efficiency; and controls upgrades. As an example, NASA installed a half-sized chiller in the utility annex—the facility that supplies chilled water to the Launch Complex 39 area—so as to better match generation capacity with the demand and reduce losses. The agency also retrofitted lighting and lighting controls with the latest in fluorescent lamp and ballast technology. In total, these multimillion-dollar projects saved tens of millions of kilowatt-hours and the associated greenhouse emissions.

Protecting Birds and the Shuttle

During the July 2005 launch of Discovery, a vulture impacted the shuttle's External Tank. With a vulture's average weight ranging from 1.4 to 2.3 kg (3 to 5 pounds), a strike at a critical area on the shuttle could have caused catastrophic damage to the vehicle. To address this issue, NASA formed the avian abatement team. The overall goal was to increase mission safety while dispersing the vulture population at Kennedy Space Center (KSC).

Through its research, the team attributed the large vulture population to an abundant food source—carrion (road kill). A large educational awareness effort was put into place for the KSC workforce and local visitors. This effort included determining wildlife crossing hot spots, ensuring the placement of appropriate signage on the roadways to increase traveler awareness, and timely disposal of the carrion.

NASA added new radar and video imaging systems to electronically monitor and track birds at the pads. Already proven effective, the avian radar—known as Aircraft Birdstrike Avoidance Radar—provided horizontal and vertical scanning and could monitor either launch pad for the movement of vultures. If data relayed from the avian radar indicated large birds were dangerously close to the vehicle, controllers could hold the countdown.



Endeavour, STS-100 (2001), roars into space, startling a flock of birds.

In addition to the energy-saving benefits of the projects, NASA was also able to modernize KSC infrastructure and improve facility capability. As an example, when the Vertical Assembly Building transfer aisle lighting was redesigned, better local control and energy saving fixtures were provided. At the same time, this increased light levels and color rendering capability. As another example, although KSC had a 10-megawatt emergency generator plant capable of servicing critical loads in a power outage, this same plant could not start the chillers needed for cooling these systems. As such, the backup plant was unable to sustain these loads for more than a few minutes before overheating conditions began. Soft start drives were installed on two

of the five chiller motors, thus allowing the motors to be started from the generator plant and providing a true backup capability for the Launch Complex 39 area.

In yet another partnership with Florida Power & Light Company, KSC opened a 10-megawatt solar power plant on 24 hectares (60 acres) of old citrus groves. This plant could generate enough electricity for more than 1,000 homes and reduce annual carbon dioxide emissions by more than 227,000 tons. Florida Power & Light Company estimated that the 35,000 highly efficient photovoltaic panels were 50% more efficient than conventional solar panels. This solar power plant, in addition to the 1-megawatt plant, has been supplying

KSC with electricity since 2009. The opening of the 10-megawatt solar field made Florida the second-largest solar-power-producing state in the country.

Summary

Throughout the shuttle era, NASA was a conscientious steward of not only the taxpayer's dollars but also of nature and the environment. Not only was the space agency aware of the dangers that wildlife could pose to the shuttle, it was also aware of the dangers that humans pose to the environment and all its inhabitants. As NASA moves forward, the agency continues to take proactive steps to assure a safe and efficient coexistence.

